

Teacher Quality Policy When Supply Matters[†]

By JESSE ROTHSTEIN*

Teacher contracts that condition pay and retention on demonstrated performance can improve selection into and out of teaching. I study alternative contracts in a simulated teacher labor market that incorporates dynamic self-selection and Bayesian learning. Bonus policies create only modest incentives and thus have small effects on selection. Reductions in tenure rates can have larger effects, but must be accompanied by substantial salary increases; elimination of tenure confers little additional benefit unless firing rates are extremely high. Benefits of both bonus and tenure policies exceed costs, though optimal policies are sensitive to labor market parameters about which little is known. (JEL I21, J22, J23, J24, J31, J41, J45)

In a 2010 manifesto, 16 big-city school superintendents stated confidently that “the single most important factor determining whether students succeed in school ... is the quality of their teacher” (Klein et al. 2010, p. B01). Influential advocates promise that policies aimed at improving teacher quality can “turn our schools around” (Gates 2011). As Secretary of Education Arne Duncan puts it, “[w]e have to reward excellence. ... We also have to make it easier to get rid of teachers when learning isn’t happening” (Hiatt 2009, p. A17). And a California judge recently invalidated that state’s teacher tenure law, finding that it “impose[s] a real and appreciable impact on students’ fundamental right to equality of education.”^{1,2}

A large recent literature focuses on the measurement of teacher effectiveness (e.g., Bill and Melinda Gates Foundation 2012; Chetty, Friedman, and Rockoff 2014a). Relatively little attention has been paid to the design of policies that will use the new measures to improve educational outcomes. This is an important omission; it is not clear what these policies should look like nor how effective they are likely to be. What evidence exists is discouraging. Several recent experiments have examined

* Goldman School of Public Policy and Department of Economics, University of California, Berkeley, 2607 Hearst Avenue #7320, Berkeley, California 94720-7320 (e-mail: rothstein@berkeley.edu). I thank Sarena Goodman for excellent research assistance and David Card, Sean Corcoran, David Figlio, Richard Rothstein, Cecilia Rouse, Todd Sorensen, Chris Taber, and numerous conference and seminar participants for helpful discussions. I am grateful to three anonymous referees for their suggestions and to the Institute for Research on Labor and Employment at UC Berkeley for research funding.

[†] Go to <http://dx.doi.org/10.1257/aer.20121242> to visit the article page for additional materials and author disclosure statement(s).

¹ *Vergara v. California*, 58 Cal. BC 484642 (2014, p. 8).

² Performance bonuses and firing are not the only potential routes to improved teacher effectiveness. Alternatives include improved selection at hiring or more or better professional development. Researchers have not identified characteristics observable at the time of hiring that are strongly correlated with subsequent effectiveness (Hanushek and Rivkin 2006; Clotfelter, Ladd, and Vigdor 2007; Rockoff et al. 2011). However, Taylor and Tyler’s (2012) examination of formative evaluations of experienced teachers found large impacts on teachers’ subsequent performance.

the short-term effects of performance bonus policies in US schools, with generally disappointing results (Goodman and Turner 2013; Fryer 2013; Springer et al. 2010; though see Fryer et al. 2012).³

Many observers believe that variation in teacher effectiveness primarily reflects personality traits.⁴ Under this view, selection is the most likely route to improved instructional quality. A well-designed contract could make the profession more attractive to effective teachers and less attractive (or unavailable) to ineffective teachers (Lazear 2003).

This type of effect is difficult to study empirically. Career decisions depend on expected compensation many years in the future, and short-term experimental interventions cannot much affect this. Even quasi-experimental approaches are not promising. Performance pay programs have generally been short-lived (Murnane and Cohen 1986), so potential teachers are unlikely to expect that recent pilot programs will last very long.

This paper examines the selection effects of alternative teacher contracts. I develop a stylized model of the teacher labor market that incorporates heterogeneity in teaching ability, dynamic learning, and contracts that condition pay or retention on realized performance. Teacher supply responses derive from a dynamic discrete choice model in which graduates and experienced teachers choose between teaching and alternative occupations on the basis of anticipated compensation, which in turn depends on the (potential) teacher's prior information about her own ability.⁵ Decisions to enter teaching depend on risk-adjusted expected compensation over the whole career. Similarly, experienced teachers' exit decisions consider the expected teaching salary over the teacher's remaining career.

A consensus result in the recent teacher quality literature is that characteristics that might be observed before entry into teaching are at best weakly correlated with eventual effectiveness. Thus, I assume that teacher ability is fixed but unknown to either the employer or the teacher herself. Compensation and retention decisions can condition only on a sequence of noisy performance signals, which might be "value-added" scores or some alternative. A prospective teacher starts with a private signal about her own ability, then updates her estimate with each performance measure. A teacher who receives positive signals raises her estimate of her own ability and thus raises her subjective expectation of the number of performance bonuses she will receive in future years and lowers her estimate of the likelihood that she will be fired for poor performance, while a teacher who receives negative signals responds in the opposite way. These expectations drive the teacher's dynamic decision making about whether to enter the profession and, having entered, to remain.

Given the extremely limited variation in extant teacher contracts, I do not attempt to estimate the model. Instead, I simulate the impact of alternative contracts using

³There is more positive evidence from other countries, mostly poorer than the United States. See, e.g., Lavy (2002); Muralidharan and Sundararaman (2011); and Duflo, Hanna, and Ryan (2012); but also see Glewwe, Ilias, and Kremer (2010).

⁴Klein et al. (2010, p. B01), for example, urge us to "stop pretending that everyone who goes into the classroom has the ability and temperament" to be an effective teacher.

⁵Dynamic occupation choice models include Adda et al. (2013) and Keane and Wolpin (1997). Essentially static models of teacher attrition include Murnane and Olsen (1989) and Dolton and van der Klaauw (1999). Tincani (2011) develops a static model of teacher sectoral choice under policies like those considered here. See also Wiswall (2007).

plausible parameters, exploring the robustness of the results through extensive sensitivity checks.

My policy analysis is closely related to the personnel economics literature on incentive contracts (e.g., Prendergast 1999; Lazear 2000) and to studies of teacher firing policies (by, e.g., Staiger and Rockoff 2010; Boyd et al. 2011; and Winters and Cowen 2013a,b). The latter studies ignore teachers' behavioral responses. In Staiger and Rockoff's (2010) simulation of tenure policies, for example, teachers can be replaced with new hires without limit, with no consequence for the quality of applicants or the salary that must be paid. Not surprisingly, then, it is optimal to deny tenure to most teachers—80 percent or more in the authors' preferred specifications. My model adds a nondegenerate teacher labor market. An increased firing rate requires higher salaries, both to compensate prospective teachers for the increased risk and to attract the needed additional applicants.⁶ I find that the required salary increase is substantial. Optimal firing rates are much lower than those obtained by Staiger and Rockoff (2010). The resulting policies can be cost effective, but offer much smaller net benefits than have sometimes been promised.

My framework allows me to consider a broader class of contracts than do previous studies. I focus initially on a performance bonus based on recent outcomes and a one-time tenure decision, but I also explore contracts that allow for ongoing retention decisions throughout the career. I also examine possible interactions between credentialing requirements (which can be seen as a fixed cost to entering the profession) and contract terms.

In order to focus on the selection margin, I rule out any other effects of performance-sensitive teacher contracts. Effort is irrelevant, all relevant outputs are measured, and teachers cannot influence their actual or measured performance. This neglects the likelihood that high-stakes accountability could lead to distortion of the performance measure (Campbell 1979). High-stakes contracts can be counterproductive in this case (Baker 1992, 2002; Holmström and Milgrom 1991). I return to this topic in Section V.

I. What are the Policies of Interest?

In a three-year random assignment study of performance bonuses, Springer et al. (2010) found no effect on student outcomes. But a three-year experiment can identify only effects operating through teacher effort. Identifying selection effects experimentally would require that the researcher “start by identifying a couple thousand high school students, follow them for fifteen or twenty years, and study whether alterations to the compensation structure of teaching impacted who entered teaching, how they fared, and how it changed their career trajectory”; even if this could be accomplished, the study “wouldn't tell us what to do today [and] wouldn't generate much in the way of findings until the 2020s” (Hess 2010). Efforts to evaluate selection effects via natural experiments face similar challenges.

⁶Although layoffs during and after the Great Recession have created considerable slack in the teacher labor market, as recently as 2007 education policymakers worried about where they would find enough qualified new teachers to replace retiring teachers (Chandler 2007; Gordon, Kane, and Staiger 2006).

This motivates my structural modeling strategy. A rich enough model can be used to simulate even the long-run effects of policies that have not yet been implemented. The simulation, of course, is conditional on the parameter values used. In principle, the parameters of a correctly specified model could be identified using data on teachers paid under the traditional “single salary schedule” that ties salaries to education and experience without regard to effectiveness. But results would be highly sensitive to functional form and distributional assumptions. I rely instead on parameter values informed by the best evidence from the literature, and present extensive sensitivity analyses that vary the parameters within plausible ranges.

The model has two primary components, a performance measurement system and a specification of teacher labor supply. A large recent literature examines the former topic, usually through “value-added” models that measure a teacher’s effectiveness based on her students’ test score growth. I use estimates from this literature to calibrate the performance measurement parameters. Nothing in the model is specific to a value-added-based system, however: it could equally well describe contracts based on more traditional, observation-based performance measures.

There is less guidance in the literature about the parameters governing the labor supply portion of the model. I discuss here several aspects of the policies of interest that bear on this portion of the model.

First, I focus on policies implemented at the state or national level rather than by individual districts. This implies that the relevant labor supply elasticities are at the occupation level, and are likely much smaller than are firm-level elasticities.⁷ Relatedly, I rule out the “dance of the lemons,” in which teachers denied tenure by one district are rehired in neighboring districts. In my model, teachers who are not retained must exit the profession.

Second, an important issue in my analysis is uncertainty about a teacher’s ability that is gradually resolved through her demonstrated performance on the job. Accordingly, I model both entry and retention decisions. Because uncertainty is greatest at the beginning of the career, entry decisions are (endogenously) insensitive to the performance component of the contract, while exit decisions become gradually more sensitive as the career goes on.

Third, occupation choices depend on the trajectory of anticipated salaries over the career. While in my model teachers’ expectations incorporate uncertainty about their own abilities and noise in the performance measure, I rule out uncertainty about the future direction of policy: prospective teachers assume the contract under consideration will be in effect throughout their careers, and I examine steady-state effects after all teachers recruited under a prior contract have retired.

Finally, teacher quality depends on both supply and demand. I assume districts are unable to distinguish teacher ability at the point of hiring.⁸ Their only options are to

⁷Lazear’s (2000) Safelite Auto Glass study examines a firm-level performance pay program; a similar program implemented at the industry level would likely have smaller selection effects. The ongoing evaluation of the federal Teacher Incentive Fund will assign schools to treatments within participating districts (Glazerman et al. 2011), so at best will identify the partial equilibrium effects of locally implemented policies.

⁸Ballou (1996) suggests that quality is not rewarded in teacher hiring decisions. Rockoff et al. (2011) find that information available at the time of hiring is only weakly predictive of effectiveness. My assumption implies that across-the-board salary changes have no effect on quality (though see Figlio 2002). This is not inconsistent with a long-run decline in teacher quality as high ability women’s nonteaching options improved (see, e.g., Corcoran, Evans, and Schwab 2004), as the latter trend had differential effects on relative pay by ability.

adjust contract parameters to induce self-selection on the part of potential teachers and/or to retain experienced teachers selectively based on observed performance.

II. The Model

I develop the model in three parts. Section IIA defines the performance measure and the Bayesian learning process, Section IIB introduces the on-the-job search model that governs entry and exit decisions, and Section IIC describes the performance-linked contracts.

A. Effectiveness, Performance Measurement, and Learning

Individual i has fixed ability τ_i as a teacher. In the current pool of teachers, ability is normally distributed with mean 0 and standard deviation σ_τ , though new contracts may change the selection process and thereby alter that distribution.

A teacher's output (per student taught) depends on her ability; her experience, t , with return to experience $r(t)$; and the size of her class, c : $y_{it}^* = \tau_i + r(t) + \gamma \ln(c)$. Each year, a noisy productivity measure is observed by both the teacher and the employer:⁹

$$(1) \quad y_{it} = \tau_i + \epsilon_{it}.$$

The noise component, ϵ , is i.i.d. Gaussian with mean 0 and standard deviation σ_ϵ . The performance measure is unbiased: all teachers draw their ϵ s from the same distribution, regardless of who they teach or the methods they use.¹⁰

Prospective teachers have limited information about their τ_i s. At entry, teacher i 's prior is $\tau_i \sim \mathcal{N}(\mu_i, \sigma_\tau^2 - \sigma_\mu^2)$, where μ_i represents the teacher's private information and $\mu_i \sim \mathcal{N}(0, \sigma_\mu^2)$ in the population of current teachers. The precision of potential teachers' information can be summarized by

$$(2) \quad h \equiv \frac{V(E[\tau | \mu])}{V(\tau)} = \frac{\sigma_\mu^2}{\sigma_\tau^2},$$

where $h = 1$ corresponds to perfect accuracy and $h = 0$ to a total lack of information. The employer observes neither μ_i or τ_i , so can base compensation and retention only on the y_{it} sequence.

Incumbent teachers update their priors rationally as performance signals arrive. The teacher's posterior after t years is

$$(3) \quad \tau | \boldsymbol{\theta}_t \sim \mathcal{N}\left(\frac{t^{-1}\sigma_\epsilon^2\mu + (1-h)\sigma_\tau^2\bar{y}_t}{t^{-1}\sigma_\epsilon^2 + (1-h)\sigma_\tau^2}, \frac{1}{t\sigma_\epsilon^{-2} + (1-h)^{-1}\sigma_\tau^{-2}}\right),$$

⁹In practice, the signal is $\tau_i + r(t) + \gamma \ln(c) + \epsilon_{it}$. But t , c , γ , and the $r(\cdot)$ function are known.

¹⁰Chetty, Friedman, and Rockoff (2014a) and Kane et al. (2013) argue that real-world performance measures have this property, though see Rothstein (2010, 2014) and Rothstein and Mathis (2013).

where $\theta_t \equiv \{\mu, y_1, \dots, y_t\}$ and $\bar{y}_t \equiv t^{-1} \sum_{s=1}^t y_s$ is the average performance signal to date. I denote the teacher's posterior mean—the first term in (3)—by $\hat{\tau}_{it}$. $\hat{\tau}_{i0} = \mu_i$, but as t grows the influence of the original guess shrinks and $\hat{\tau}_{it}$ converges toward τ_i .

B. The Teacher Labor Market: Entry and Persistence

Prospective teachers have von Neumann-Morgenstern utility $u(w)$, defined over annual compensation w , and discount rate δ . A prospective teacher with information μ draws a single nonteaching job offer, providing continuation value ω_1 , from a distribution Ω_1 . She compares this to the utility she will obtain from a teaching career. I denote this by $V_1(\mu; C)$ to emphasize that it may depend both on μ and on the terms of the contract C . She enters teaching if $V_1(\mu; C) > \omega_1$.

Each year of teaching represents a new stage of the dynamic decision game. A teacher beginning her t th year, $1 \leq t \leq T$, has information θ_{t-1} . In year t , she receives a performance signal y_t and is paid a salary w_t . This salary may depend on past and/or current performance. The employer then decides whether to offer her continued employment in $t + 1$, again considering her performance to date. Let $f_t = f_t(y_1, \dots, y_t; C)$ be an indicator for being fired after period t , and let V_{t+1}^f be the continuation value of a teacher who is fired after t . A teacher who is not fired updates her estimate of her own ability based on θ_t and uses this to forecast her future inside earnings and retention probability. She then draws a single outside wage offer, summarized by its continuation value $\omega_{t+1} \sim \Omega_{t+1}$, and decides whether to remain in teaching in $t + 1$ or to accept the outside offer.

The value of remaining in teaching in year $t + 1$ is $V_{t+1}(\theta_t; C)$. A teacher whose outside offer ω_{t+1} exceeds this value accepts the offer. Teachers who accept outside offers, either initially or later, can not reenter teaching. V_t is thus defined recursively:

$$(4) \quad V_t(\theta_{t-1}; C) = E \left[u(w_t) + \delta \left\{ f_t V_{t+1}^f + (1 - f_t) \max(\omega_{t+1}, V_{t+1}(\theta_t; C)) \right\} \middle| \theta_{t-1} \right].$$

The expectation is taken over the teacher's posterior τ distribution following period $t - 1$, as given by (3), and over the distribution of the noise term ϵ_{it} . Careers end after T periods, so $V_T(\theta_{T-1}; C) = E[u(w_T) | \theta_{T-1}]$.

The Ω_{t+1} distribution ($t < T$) is calibrated so that the annual exit hazard under the base contract C_0 (discussed below) is λ_0 and the elasticities of entry and exit with respect to certain, permanent changes in w are η and $-\zeta$, respectively.¹¹ The online Appendix discusses the censored Pareto distribution that generates this.

ω_{t+1} is assumed independent of θ_t and τ . The available evidence indicates little relationship between teaching effectiveness and traditional human capital measures (Rockoff et al. 2011). Several studies find negative correlations between effectiveness and exit from teaching (Krieg 2006; Goldhaber, Gross, and Player 2011). Given the weak or nonexistent pecuniary returns to effectiveness in teaching, one

¹¹ As $T \rightarrow \infty$ the average career length approaches λ_0^{-1} and the elasticity of the career length with respect to the inside wage converges to ζ . With the parameters used below ($T = 30$, $\lambda_0 = 0.08$, and $\zeta = 1$), the average career length is 11.5 years (versus 12.5 years with $T = \infty$) and the career length elasticity is roughly $0.77\zeta = 0.77$. The total labor supply elasticity is the sum of the entry and career length elasticities, approximately $\eta + 0.77\zeta = 1.77$.

would expect the opposite if teaching ability were positively correlated with outside wages.¹² Nevertheless, I weaken this assumption later.

An important parameter governing the effect of tenure contracts is V_{t+1}^f , the continuation value of a teacher who is fired. I assume $V_{t+1}^f = (1 - \kappa_t) V_{t+1}(\theta_t; C_0)$, where $0 < \kappa_t < 1$ represents the penalty for being fired after t relative to being retained under contract C_0 . (Note that with my assumptions, $V_{t+1}(\theta_t; C_0)$ is invariant to θ_t .) A teacher who exits without being fired does not pay the penalty. Thus, if κ_t is large, teachers who expect to be fired with high probability will instead exit voluntarily beforehand.

I evaluate V_t numerically, using an algorithm described in the online Appendix. To simulate the impact of alternative contract C , I draw teachers from the $\{\mu, \tau\}$ distribution, then draw performance measures $\{y_1, \dots, y_T\}$ for each. For each teacher and each year t , I compute $V_t(\theta_{t-1}; C)$ and $V_t(\theta_{t-1}; C_0)$. The ratio of these, along with the Ω_t distribution, determines the probability of entering the profession and, conditional on entering, of remaining from $t - 1$ to t . Note that my assumption of a constant labor supply elasticity allows me to avoid modeling explicitly the distribution of μ in the *population* of potential teachers—changes in the returns to teaching induce proportional changes in the amount of labor supplied by each μ type.

C. Teacher Contracts

The baseline contract C_0 ties pay to experience but not to performance: $w_{it}^0 = w^0(1 + g(t))$, with $g'(\cdot) \geq 0$. No teachers are fired: $f_{it} \equiv 0$. Alternative contracts base either w_{it} or retention decisions on the sequence of performance signals to date. I consider two alternatives, performance-based bonuses and performance-based tenure decisions:

Bonuses.—Bonuses are awarded to teachers with high measured performance, averaged across two years to reduce the influence of noise. Thus, in year t all teachers with $\frac{y_{it} + y_{i,t-1}}{2} \geq y^B$ receive bonuses; first-year teachers are ineligible. Total compensation is $w_{it}^B = \alpha^B w_{it}^0(1 + b \times e_{it})$, where e_{it} is an indicator for bonus receipt, b indexes the size of the bonus (as a share of base pay), and α^B is an adjustment to base pay relative to the baseline contract. The threshold y^B is set to ensure that in the absence of behavioral responses a share s^B of teachers would receive bonuses each year.

Tenure.—Teachers are evaluated for tenure after their second years. Any teacher whose average performance to date $\frac{y_{i1} + y_{i2}}{2}$ exceeds a threshold y^T is given security of employment. Those falling short of the threshold are dismissed. y^T is calibrated so that a share s^T of current entrants are tenured. As before, this threshold is fixed; if

¹²Chingos and West (2012) find that former teachers' salaries are positively correlated with their value-added as teachers. But they also find that value-added is uncorrelated with attrition rates (West and Chingos 2009). One potential explanation is that τ is positively correlated both with outside salaries and with the individual's taste for teaching, leaving no correlation between τ and the net desirability of a nonteaching offer (at least among those who currently select into the profession). It is the latter that is relevant here.

the ability distribution of new recruits rises, so will the tenure rate. Pay is as under C_0 with an adjustment α^T : $w_{it}^T = \alpha^T w_{it}^0$.

In my model, the optimal pay schedule would have low annual pay and a very large retirement bonus that depends on performance throughout the career. This is unrealistic, but it is not obviously unreasonable to incorporate information after year 2 into retention decisions. In Section IVB I explore the optimal choices of s^B , b , and s^T under the above contract structures, as well as alternative firing rules that better use the available information.

I consider two scenarios for the choices of α^B and α^T . First, I assume that demand is inelastic; base wages are set (via α^B and α^T) to yield the same number of teachers as under C_0 . This is consistent with laws and collective bargaining contracts that commonly specify class sizes. I use this scenario to explore the effect of alternative contracts on total compensation costs. Second, I assume instead that the district's budget is fixed, so that increases in teacher salaries must be offset via reductions in the number of teachers hired and thus via increases in class size. Here, the α parameters are set so that the resulting total labor supply (not counting that which teachers who have been fired would like to supply) exhausts the baseline budget at the specified salaries, perhaps with more or fewer classrooms than under C_0 . This is more consistent with a budgeting process that treats revenues as exogenous and balances the budget via adjustments to the workforce size. I assume that balance is achieved over the long run, so that changes in the number of teachers are matched by equivalent savings on facilities and all other expenses.¹³ The fixed budget scenario allows me to explore the cost effectiveness of alternative contracts relative to traditional uses of school resources.

D. Parameter Values

My primary simulations set parameters at what I judge to be likely values, attempting to err on the side of optimism about the prospects for performance contracts. These parameters are shown in column 1 of Table 1. For parameters for which I have less evidence, I also consider a pessimistic scenario (column 2), in which I expect the performance contracts to be less effective, and an optimistic scenario (column 3).

I calibrate the productivity and performance measurement parameters using value-added literature. The standard deviation (SD) of teacher value-added for students' end-of-year test scores has been widely estimated to be between 0.1 and 0.2, with 0.15 as a reasonable central estimate (e.g., Rivkin, Hanushek, and Kain 2005; Rothstein 2010; Chetty, Friedman, and Rockoff 2014a). Many studies have found that teachers improve with experience but level off quickly; I draw $r(t)$ from Staiger and Rockoff's (2010) estimates for New York City. When I allow the district to vary class sizes to offset changes in teacher salaries, I assume a 1 percent increase in class size reduces student achievement by 0.004 standard deviations. This is based on the STAR class size experiment, in which students in small classes, averaging

¹³ If salary costs were the only component of the budget, my fixed budget assumption would correspond to a unit labor demand elasticity. Non-salary costs reduce the absolute demand elasticity below one.

TABLE 1—VALUES OF KEY PARAMETERS

Parameter	Scenario			
	Baseline (1)	Pessimistic (2)	Optimistic (3)	
<i>Effectiveness</i>				
σ_τ	SD of teacher ability	0.15		
$r(t)$	Experience effect on productivity	$\left\{ \begin{array}{l} -0.07 \text{ if } t = 0 \\ -0.04 \text{ if } t = 1 \\ -0.02 \text{ if } t = 2 \\ 0 \text{ if } t > 2 \end{array} \right.$		
γ	Effect of 1% increase in class size		-0.004	
<i>Measurement</i>				
σ_ε	SD of noise in annual performance		0.18	
<i>Teacher preferences and information</i>				
h	Reliability of private information	0.1	0.05	0.25
δ	Discount rate		0.97	
$u()$	von Neumann-Morganstern utility	Linear	CRRRA, $\rho = 3$	Linear
η	Elasticity of entry w.r.t. w_0	1	0.5	1.5
ζ	Absolute elasticity of exit w.r.t. w_0	1	0.5	1.5
λ_0	Annual exit hazard, base contract		0.08	
T	Maximum career length		30	
κ	Effect of being fired	1%* min(t , 10%)	15%	1%* min(t , 5%)
<i>Base contract</i>				
$g(t)$	Return to experience		0.015* t	
<i>Bonus contract</i>				
b	Bonus size (as share of base pay)		20%	
s^B	Share of teachers receiving bonus		25%	
α^B	Base pay adjustment (calibrated)			
	Fixed quantity scenario	-3.5%	-2.7%	-3.5%
	Fixed budget scenario	-3.5%	-2.9%	-3.5%
<i>Tenure contract</i>				
s^T	Share of teachers receiving tenure		80%	
α^T	Base pay adjustment (calibrated)			
	Fixed quantity scenario	+12.6%	+57.8%	+8.5%
	Fixed budget scenario	+10.4%	+27.1%	+7.4%

15 students, outperformed those in large classes, averaging 22, by 0.15 standard deviations (Krueger 1999).

I set $\sigma_\varepsilon = 0.183$. The reliability of y (defined as $\frac{V(\tau)}{V(y)} = \frac{\sigma_\tau^2}{\sigma_\tau^2 + \sigma_\varepsilon^2}$) is then 0.4, near the top of the range identified in Sass' (2008) survey of value-added measurement.

The parameter h quantifies the information that prospective teachers have about their own abilities. When Springer et al. (2010) asked experienced teachers (with several years of performance measures under their belts) to forecast their probabilities of winning performance awards, their forecasts were uncorrelated with actual award receipt. This is inconsistent with a large h . Moreover, several studies find that observable teacher characteristics are poor predictors of future effectiveness. The strongest correlations come from Rockoff et al. (2011), who find that a rich vector of academic and personality characteristics explains only 10 percent of the variance in value-added. This corresponds to $h = 0.1$.

One might expect teachers—who have self-selected into a very secure but low paying occupation—to be unusually risk averse (Flyer and Rosen 1997). I use linear utility as a baseline, but also consider in my pessimistic parameters constant relative risk aversion (CRRA), with coefficient 3, over annual pay.¹⁴ I use a 3 percent discount rate in all scenarios.

The outside salary offer distribution is set to yield a constant 8 percent annual exit hazard under contract C_0 . Careers end after $T = 30$ years. This is roughly consistent with the observed national data, though in these data exit rates are somewhat higher in the first years of teachers' careers; see online Appendix Figure 1.

Ransom and Sims (2010) examine how teacher retention probabilities vary with the district's wage schedule. Their estimates imply $\zeta = 1.8$. Ransom and Sims interpret their estimates in a monopsony framework, and many of the teachers in their sample who quit may take jobs in neighboring districts. Their estimate thus likely overstates the occupation-level exit elasticity, though it is not clear by how much.¹⁵ I assume that $\zeta = 1$ at the occupation-level, but also consider $\zeta = 0.5$ and $\zeta = 1.5$. Following Manning (2005), I assume the entry elasticity equals the exit elasticity: $|\eta| = |\zeta|$.

κ represents the permanent effect on future earnings of being fired for poor performance. I assume that a teacher denied tenure in year 2 sees her future earnings reduced by 2 percent; when I vary the date of firing decisions, in Section IVB, a teacher fired after t years suffers a $\min\{t, 10\}$ percent reduction. This choice of κ is more likely to be too small than too large. By my calculations, Dee and Wyckoff's (2013) study of performance-based firing threats under the Washington, DC IMPACT program implies that κ is one or more orders of magnitude larger than this.¹⁶

I assume that base (real) teaching pay rises by 1.5 percent with each year of experience.¹⁷ The bonus contract provides a $b = 20$ percent bonus for teachers whose two-year moving average performance exceeds a fixed threshold $y^B = +0.178$, set to ensure that $s^B = 25$ percent of the current teaching workforce would get bonuses. The tenure contract is calibrated to yield a tenure rate of $s^T = 80$ percent given the current ability distribution, corresponding to a threshold of $y^T = -0.167$. Both y^B and y^T are fixed: if the alternative contracts attract more high- τ teachers then more bonuses would be paid or more teachers would be tenured.

¹⁴To my knowledge, mine is the first dynamic occupation choice model to allow for risk aversion. A more complete model would define $u(\cdot)$ over annual consumption and allow agents to borrow and save. This would make bonus contracts more attractive but would also add considerable complexity. Rothstein (2012) presents an alternative model of risk aversion that captures some income smoothing.

¹⁵Clotfelter et al. (2008) study a targeted (but not performance-dependent) bonus program and estimate a school-level exit elasticity between 3 and 4. They do not distinguish exits from the profession, movements to other districts, and movements to other schools in the same district.

¹⁶Under IMPACT, teachers who receive two consecutive *minimally effective* (ME) ratings face dismissal. Dee and Wyckoff (2013), using a regression discontinuity design, find that an initial ME rating increases the annual exit rate by over one-third. As only about 14 percent of teachers near the ME threshold in one year receive an ME the next year (personal communication from Thomas Dee), the implied value of $\zeta\kappa$ is 2.6. Dee and Wyckoff's (2013) parallel analysis of increases in the likelihood of future performance-linked pay increases yields a point estimate of $\zeta = 2.6$, though with a very wide confidence interval. Combining these implies $\kappa = 100$ percent; lower estimates of ζ yield even larger κ . In a quite different setting, von Wachter, Song, and Manchester (2009) find that displacement in mass layoffs reduces earnings by 20–30 percent, with effects that persist for at least 20 years. Laid-off workers were often older and displaced from declining occupations and industries; on the other hand, I assume that *all* fired teachers must move to new occupations and industries.

¹⁷Teacher pay is often back-loaded, particularly when pension accumulations are counted. This may serve to lock in midcareer teachers, though empirical exit hazards (see online Appendix Figure 1) are nontrivial at all t .

The final parameters are α^B and α^T , the adjustments to base pay under the bonus and tenure contracts. These are calibrated, given the other parameters, either to ensure the same total number of teachers (in steady state) as are obtained under the baseline contract or to satisfy a fixed budget constraint. In the latter calibrations, I assume that a 1 percent reduction in the number of teachers (i.e., a 1 percent increase in class size) would produce savings equal to 3 percent of the average teacher's salary.¹⁸ Under my baseline parameters, the bonus contract requires a 3.5 percent reduction in base pay under either demand scenario; the tenure contract requires base salaries to increase by 12.6 percent to maintain the same number of teachers or by 10.4 percent to balance a fixed budget. The pessimistic parameters imply higher salaries under each contract and demand scenario, while salaries under the tenure contract are somewhat lower with the optimistic parameters.

III. Results

A. Noise, Information, and Incentives

The incentive faced by a teacher i with prior $\hat{\tau}_{it}$ depends on the link between this prior and her true ability, the link from ability to the performance signal, and the link from the signal to the contract terms. Moreover, the success of a contract depends on the average incentive perceived by teachers at each true ability level τ_i , among whom there may be much variation in $\hat{\tau}_{it}$. Each of these links serves to dampen the incentives for self-selection.

This is easiest to illustrate for the bonus contract. By iterated expectations, the average subjective probability as of year t of receiving a bonus in year $t' \geq t + 2$ (so that year- t performance does not enter directly) among teachers of true ability τ can be expressed as:

$$(5) \quad E\left[E\left[e_{it'} \mid \hat{\tau}_{it}\right] \mid \tau_i = \tau\right] = E\left[E\left[E\left[E\left[e_{it'} \mid y_{it'}, y_{i,t'-1}\right] \mid \tau_i\right] \mid \hat{\tau}_{it}\right] \mid \tau_i = \tau\right].$$

The outer conditioning variables can be omitted from inner expectations because the inner variables capture all relevant information: bonus receipt is independent of ability conditional on measured performance, performance depends only on true ability and not on subjective perceptions, and these perceptions depend only on true ability only through $\hat{\tau}_{it}$.

The innermost expectation on the right side of (5) is a step function, as $e_{it'} \equiv \mathbf{1}\left(\frac{y_{it'} + y_{i,t'-1}}{2} \geq y^B\right)$. But each of the three outer expectations serves to smooth this out.

First (working our way outward) consider $E\left[e_{it'} \mid \tau_i\right] = E\left[E\left[e_{it'} \mid y_{it'}, y_{i,t'-1}\right] \mid \tau_i\right]$. Using the parameters from Table 1, teachers with τ at the ninetieth percentile win bonuses only 54 percent of the time, while those at the fiftieth percentile do so 9 percent of the time.¹⁹

¹⁸Teacher salaries represent about one-third of total educational expenditures; I assume all other costs are variable in the long term.

¹⁹I express ability in terms of the percentile rank within the *baseline* τ distribution. Of course, this distribution would change under alternative contracts. The fixed-norm percentiles are simply a convenient scale.

This is smoothed out further by teachers' uncertainty about their own abilities. At every t , $V[\tau_i | \hat{\tau}_{it}] > 0$. This flattens the $E[e_{it} | \hat{\tau}_{it}] = E[E[e_{it} | \tau_i] | \hat{\tau}_{it}]$ function; even teachers who think they are likely to be of low ability realize that they might in fact be of higher ability and thus be contenders for bonuses, while those who think they are of high ability harbor doubts about this. This is particularly true of early career teachers, for whom $V[\tau_i | \hat{\tau}_{it}]$ is large. Even a prospective teacher at the ninetieth percentile of the μ distribution thinks she has only a 37 percent chance of receiving a bonus in any given year of her career, while a prospective teacher with tenth percentile μ anticipates a 4 percent chance. As teachers accumulate information, they quickly learn their places in the distribution. After one year, the teacher at the ninetieth percentile of the $\hat{\tau}_{i1}$ distribution thinks her chance of receiving a bonus is 42 percent, and this rises to 45 percent after two years and 48 percent after five years.

The $E[e_{it} | \hat{\tau}_{it}]$ curve governs the incentives faced by teachers of different $\hat{\tau}_{it}$ s. But it does no good to attract low ability teachers who think they are high ability, nor to repel high ability teachers who underestimate their own τ s. The degree to which the contract attracts teachers who are actually of high ability depends on $E[E[e_{it} | \hat{\tau}_{it}] | \tau_i]$. This further attenuates the incentives, again more so early in the career: At entry, the average ninetieth percentile teacher's subjective expectation of her own ability puts her at only the seventieth percentile.

The solid line in panel A of Figure 1 shows the average probability of winning a bonus by *true* ability percentile, $E[e_{it} | \tau_i = \tau]$. The other series in this panel show average *subjective* expectations among teachers of each true ability, $E[E[e_{it} | \hat{\tau}_{it}] | \tau_i = \tau]$, at several points in the career. On entering teaching there is relatively little differentiation except at the extreme tails of the distribution. The average ninetieth percentile teacher anticipates a 30 percent of earning a bonus in any given year while the average tenth percentile teacher perceives a 9 percent chance. But perceived incentives become much better targeted with experience. After five years, ninetieth percentile teachers perceive their chances at 44 percent, on average, while tenth percentile teachers see theirs as under 2 percent. Thus, while incentive effects of a bonus system are weak at the recruitment stage, later attrition decisions may be more sensitive.

Panel B of Figure 1 repeats the exercise for the tenure contract. (I omit the curve for fifth-year teachers, as tenure decisions have been made by then.) Again, we see weak incentives for low ability potential teachers to select other careers at the entry point, but after even a single year the incentives are stronger.

There is a close, albeit imperfect, mapping from the subjective probabilities of positive and negative outcomes graphed in Figure 1 to the average values of teachers of different abilities under the two contracts. Figure 2 shows average continuation values V of teachers under the two contracts, by ability level and years of experience.²⁰ Because the V scale is not intuitive, I convert it to equivalent variations; Changes in w^0 that would yield the same values under the single salary contract.

²⁰ Averages are computed over all teachers who enter under the baseline contract, ignoring voluntary exits. Note that the probabilities in Figure 1 depend only on the performance measurement parameters, while those in Figure 2 depend also on the labor supply, labor demand, and outside wage offer parameters. I use the *fixed quantity* (inelastic labor demand) α parameters in Figures 2–4.

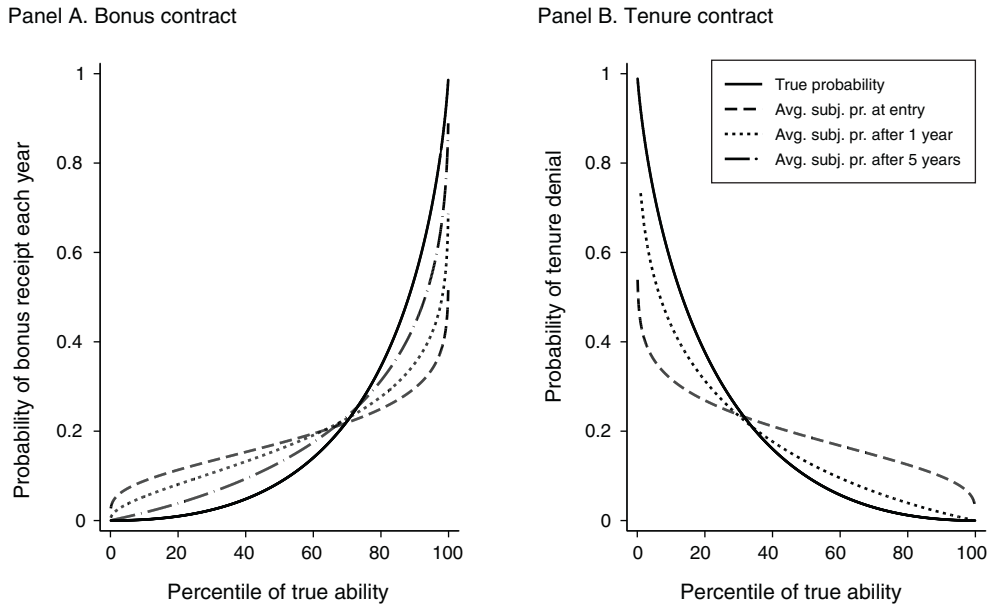


FIGURE 1. OBJECTIVE PROBABILITIES OF BONUS RECEIPT AND TENURE DENIAL AND AVERAGE SUBJECTIVE EXPECTATIONS, BY TRUE ABILITY PERCENTILE AND EXPERIENCE

Notes: Panel A shows mean objective (solid line) and subjective (dotted and dashed lines) probabilities of receiving a bonus in any year two or more years in the future. Panel B shows objective and subjective probabilities of being denied tenure after year two. All probabilities are averaged over all teachers with true ability τ , ignoring any labor supply responses to the alternative contracts. Ability is rescaled to a percentile score (in the baseline distribution), on the x -axis.

An estimate of +5 percent means that w^0 would need to rise by 5 percent under contract C_0 to yield the same value as is obtained under the alternative contract. The figure shows that the bonus contract produces the equivalent of a 0.3 percent salary increase for the average ninetieth percentile teacher at entry, and a 4.9 percent increase after five years.²¹ These changes are not large, and suggest that any self-selection responses to the bonus contract will be quite modest. The tenure contract achieves a much steeper slope among first-year teachers, with a range of about 10 percent of baseline salaries between tenth and ninetieth percentile teachers, but like the bonus contract does little to alter the relative incentives governing entry to teaching. After the tenure decision teachers benefit from the increased salaries needed to attract sufficient applicants, and values reflect the 12.6 percent increase in salaries across all ability levels.

B. Impact of Incentives

I next turn to the impacts of the contracts on the teacher ability distribution. Panel A of Figure 3 shows how the two contracts would change the number of entering teachers at each ability percentile, relative to the baseline contract. The bonus

²¹ Recall that base salaries are adjusted by a factor α^B under this contract. The equivalent variation is thus $\alpha^B - 1 = -3.5$ percent for a teacher with zero subjective probability of future bonus receipt.

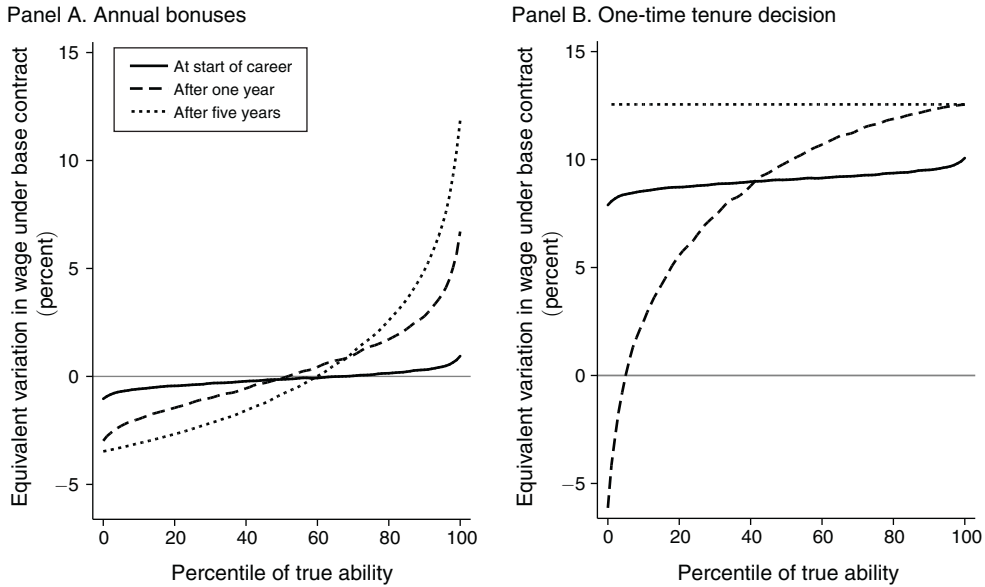


FIGURE 2. AVERAGE EQUIVALENT VARIATION OF BONUS AND TENURE CONTRACTS, BY ABILITY AND EXPERIENCE

Notes: Series show the changes in base wage (w^0) under the baseline contract that would be needed to match the average value obtained by teachers at different points in the ability-experience ($\tau - t$) distribution under the performance-based contracts with wages set to fix the number of teachers employed. Ability is rescaled to a percentile score (in the baseline distribution), on the x -axis. For example, the point (0.95, 3.8) on the dashed line in panel A indicates that under the bonus contract the average teacher with τ at the ninety-fifth percentile of the baseline distribution and one year of experience obtains a value equivalent to what would be obtained with a 3.8 percent salary increase under the single salary contract.




FIGURE 3. EFFECT OF ALTERNATIVE CONTRACTS ON RECRUITMENT AND CAREER LENGTH, BY ABILITY PERCENTILE

Notes: Figure shows percentage change, relative to the single salary contract, in the number of new hires (panel A) and the average career length (panel B) at each ability (τ) level under each of the performance-based contracts, when wages are set to fix the number of teachers employed. Ability is rescaled to a percentile score in the baseline distribution, on the x -axis.




FIGURE 4. EFFECT OF ALTERNATIVE CONTRACTS ON TOTAL NUMBER OF TEACHERS, BY ABILITY PERCENTILE

Notes: Figure shows percentage change, relative to the single salary contract, in the total number of teachers employed at each ability (τ) level under each of the performance-based contracts, when wages are set to fix the overall workforce size. Ability is rescaled to a percentile score (in the baseline distribution), on the x -axis.

contract entices more high ability and fewer low ability entrants, but the impact is extremely small. The tenure contract, with its sharply increased salaries, attracts more of all types, but again does little to change the relative numbers of high and low ability teachers.

Panel B of Figure 3 shows the effects of the two contracts on average career length. The bonus contract has small effects on this margin, too, concentrated at the very top of the ability distribution. The $1 - \alpha^B = 3.5$ percent cut to baseline salaries under this contract reduces career lengths of below-median teachers by only about 2 percent.

The tenure contract has a more dramatic effect. Career lengths shorten by as much as 80 percent for the weakest teachers. This primarily reflects tenure denials. Among slightly higher ability teachers, around the twentieth percentile, career lengths shorten by an average of about 25 percent. This change would be less than one-third as large in the absence of labor supply responses; the bulk of it derives from voluntary exits after the first year among teachers who have learned their tenure chances are lower than anticipated but who would nevertheless receive tenure if they stayed. At the other end of the distribution, careers lengthen by nearly 10 percent, as higher salaries reduce voluntary attrition among tenured teachers.

Figure 4 presents the impact of the two contracts on the steady-state number of teachers at each ability level, combining entry and career length effects. Not surprisingly, the bonus contract has little effect, reducing the number of low ability teachers by about 3 percent and increasing the number of high ability teachers by a bit more than this. The tenure policy is much more effective, increasing the number of classes taught by top-quartile teachers by about 20 percent and reducing those taught by bottom decile teachers by 60 percent.

TABLE 2—IMPACT OF BONUS AND FIRING CONTRACTS ON TEACHER EFFECTIVENESS AND TOTAL COSTS

	Baseline (1)	Inelastic demand		Fixed budget	
		Bonuses (2)	Tenure (3)	Bonuses (4)	Tenure (5)
<i>Outcomes under alternative contracts</i>					
Number of teachers	1	1.00	1.00	1.00	0.97
Teacher ability (τ)					
Mean	0.000	0.004	0.033	0.004	0.033
SD	[0.150]	[0.151]	[0.135]	[0.151]	[0.135]
Teacher experience					
Percent 1st year	8.0	8.0	8.9	8.0	9.0
Percent 1st 3 years	30.9	30.9	31.1	30.9	31.6
Mean	8.82	8.83	9.09	8.83	8.99
Teacher effectiveness ($\tau+r(t)$)					
Mean	-0.011	-0.007	0.021	-0.007	0.021
SD	[0.151]	[0.153]	[0.139]	[0.153]	[0.139]
Class size effect (rel. to baseline)	0	0.000	0.000	0.000	-0.014
Average output (teacher effect + class size)	-0.011	-0.007	0.021	-0.007	0.007
Salaries (expressed as multiple of baseline starting salary)					
Base starting salary	1.000	0.965	1.126	0.965	1.104
Average total pay	1.148	1.147	1.298	1.147	1.271
<i>Change from baseline</i>					
Number of teachers		n/a	n/a	+0.01%	-3.4%
Teacher ability (τ)					
Mean		+0.004	+0.033	+0.004	+0.033
SD		+0.001	-0.015	+0.001	-0.015
Teacher experience					
Percent 1st year		-0.0 p.p.	+0.8 p.p.	-0.0 p.p.	+1.0 p.p.
Percent 1st 3 years		-0.0 p.p.	+0.3 p.p.	-0.0 p.p.	+0.7 p.p.
Mean		+0.009	+0.270	+0.009	+0.167
Teacher effectiveness ($\tau+r(t)$)					
Mean		+0.004	+0.033	+0.004	+0.032
SD		+0.001	-0.012	+0.001	-0.012
Class size effect		+0.000	-0.000	+0.000	-0.014
Average output (teacher effect + class size)		+0.004	+0.033	+0.004	+0.018
Salaries					
Base starting salary		-3.5%	+12.6%	-3.5%	+10.4%
Average total pay		-0.1%	+15.0%	-0.0%	+12.3%

Note: Simulations use baseline parameter values from column 1 of Table 1.

Columns 2 and 3 of Table 2 show the effects of the two contracts on teacher ability, experience, effectiveness (combining ability and experience effects), and salaries. The bonus contract yields only small increases in average teacher ability, around 2 percent of a teacher-level standard deviation, while the tenure policy's effect is over eight times this. The tenure contract increases the number of first-year teachers by about an eighth, but this has little impact on overall productivity; first-year teachers are about one-half of a standard deviation less productive than experienced teachers, so a 0.9 percentage point increase in inexperienced teachers reduces average productivity by less than one-two-hundredth of a standard deviation. Thus, net effects on teacher effectiveness shown in the lower panel (+0.004 student-level standard deviations for bonuses relative to the baseline, and +0.033 for tenure) are the same

to three digits as the gross effects on teacher ability. Total salary costs are essentially unchanged under the bonus contract but rise by 15 percent under the tenure contract.

One way to analyze the cost-benefit trade-off is to monetize the output improvements that the contracts yield. Chetty, Friedman, and Rockoff (2014b) find that one (teacher-level) standard deviation in elementary teachers' effectiveness is associated with nearly \$200,000 in present-discounted future earnings per classroom taught.²² With an average teacher salary of \$50,000, this implies a benefit-cost ratio for the tenure contract of nearly 6 to 1.²³ If the Chetty, Friedman, and Rockoff (2014b) results are correct, then, it would be worth increasing education budgets to finance the higher salaries necessitated by alternative contracts (though see Rothstein 2014).

Another approach to cost-benefit analysis recognizes that education budgets may not be set optimally. If budgets are fixed, the cost effectiveness of the teacher contracts should be compared to that of alternative uses for school funds. These alternative uses may also have positive net benefits: the Chetty et al. (2011) analysis of class size reduction implies a benefit-cost ratio around 2.5 to 1 (see also Krueger 1999).²⁴

Columns 4 and 5 of Table 2 present fixed budget analyses of the bonus and tenure contracts, assuming that increased per-teacher costs must be offset by reducing the number of teachers per student. The tenure contract requires increasing class size by 3.4 percent to finance the higher salaries that it requires. (These are lower than in column 3, as with larger classes the district needs fewer teachers.) The negative effect of larger classes on student achievement offsets a bit less than one-half of the benefit of improved teacher quality. The net effect of the policy is to raise average output by 0.018 student-level standard deviations (0.12 teacher-level SDs). The bonus contract is also cost effective relative to class size reduction, but its impact is less than one-quarter as large.

C. Sensitivity to Alternative Parameters

Table 3 presents estimates for alternative parameter values. Here and hereafter, I focus on the *fixed budget* demand scenario, as this allows me to summarize the impact of alternative contracts by the constant-budget impact on average output, incorporating class size effects. The entries in the first row of columns 1 and 4 repeat the estimates from Table 2. Subsequent rows vary the different parameters in turn, one at a time. To illustrate potential interactions among parameters, columns 2 and 5 show results for the "pessimistic" parameter values from Table 1, while columns 3 and 6 show results for the "optimistic" parameters. (Blank cells correspond to parameter values that match those used for row 1 of the same column.)

²²This calculation is based on the Chetty, Friedman, and Rockoff (2014b) estimate of the present value of the average student's future earnings, \$522,000, 1.34 percent in additional earnings at age 28 per standard deviation of teacher value-added, and the average elementary class size in their sample, 28.2.

²³The effect on output in column 3 is 0.033 student-level standard deviations, or 0.22 teacher-level standard deviations. Thus, the benefit is $0.22 \times \$200,000$, or \$44,000.

²⁴Chetty et al. (2011) find that a one-third reduction in class size, requiring 50 percent more teachers, raises the present value of students' future earnings by \$189,000 per year. As noted above, teacher salaries average around \$50,000 and are about one-third of total education costs. Thus, a back-of-the-envelope cost estimate—ignoring effects on teacher salaries—is $(0.5)(3)(\$50,000) = \$75,000$.

TABLE 3—SENSITIVITY OF RESULTS TO MODEL PARAMETERS: EFFECTS ON TEACHER OUTPUT (*in student SDs*)

	Bonuses			Tenure		
	Baseline (1)	Pessimistic alternative (2)	Optimistic alternative (3)	Baseline (4)	Pessimistic alternative (5)	Optimistic alternative (6)
1. Baseline	+0.004	+0.001	+0.007	+0.018	-0.003	+0.023
Varying private information						
2. $h = 0.05$	+0.004		+0.006	+0.018		+0.022
3. $h = 0.25$	+0.005	+0.001		+0.019	-0.000	
Less noisy performance measure						
4. $\sigma_\varepsilon = 0.12$	+0.006	+0.001	+0.010	+0.023	+0.002	+0.028
Varying the supply elasticities						
5. $\eta = 0.5$	+0.004		+0.006	+0.015		+0.019
6. $\eta = 1.5$	+0.004	+0.001		+0.020	+0.004	
7. $\zeta = 0.5$	+0.002		+0.004	+0.014		+0.019
8. $\zeta = 1.5$	+0.006	+0.004		+0.021	+0.015	
9. ζ declines with t	+0.003	+0.002	+0.005	+0.017	+0.006	+0.021
10. η rises with μ	+0.004	+0.001	+0.007	+0.018	+0.002	+0.022
Other						
11. Risk averse (CRRA = 3)	+0.003		+0.005	+0.017		+0.022
12. High firing cost ($\kappa = -0.15$)				+0.015		+0.023
13. $\lambda = 6\%$ in years 1–4; 3% thereafter	+0.003	+0.000	+0.005	+0.018	-0.005	+0.023

Notes: Row 1 uses parameter values from Table 1. Each successive row changes one parameter at a time, as indicated; entries where the alternative parameter matches that used in row 1 are suppressed. All simulations assume a fixed budget, as in columns 4–5 of Table 2. Table entries represent impacts of the alternative contracts on average teacher productivity, in student-level standard deviations relative to the baseline contract, incorporating changes in ability, in experience, and in class size. In row 9, $\zeta = 1.5 - 0.1 \times \min(t, 10)$. In row 10, $\eta = 1.01 + 0.0649 \left(\frac{\mu}{\sigma_\mu}\right)$. This approximates the elasticity that arises in a Roy model with $\text{corr}(\mu, \omega_1) = 0.1$.

Working down columns 1 and 4 of Table 3, we can see the impact of each parameter on the results. The impacts of the two policies are not very sensitive to changes in the reliability of prospective teachers' private information (rows 2–3), but would grow noticeably if the performance measure could be made more reliable (row 4).

Rows 5–8 show the effects of varying the labor supply elasticities. Both policies are more effective when labor supply is more elastic, particularly on the exit margin. The relative unimportance of the hiring elasticity reflects the fact that entering teachers have too little information about their own abilities to perceive large changes in incentives.

Row 9 shows a variant in which the exit elasticity is allowed to decline with experience, starting at 1.5 and falling to 0.5 by year 10. This is meant to capture the intuition that early career teachers may be more mobile than are later career teachers (Ransom and Sims 2010). Results are intermediate between those with permanently higher and permanently lower elasticities; qualitative conclusions are unchanged.

Row 10 shows an additional variant in which the entry elasticity is increasing with μ : $\eta = 1.01 + 0.0649 \left(\frac{\mu}{\sigma_\mu}\right)$. In a Roy model of career choice with $\text{corr}(\mu, \omega_1) > 0$, entry of higher- μ potential teachers is more sensitive to the offered wage than that of those with lower μ . The function here is approximately what would obtain with a correlation of 0.1 and a log normal outside wage distribution calibrated to yield an average elasticity of 1. Allowing for this sort of heterogeneity in entry elasticities has little impact on the results.

In row 11 of Table 3, I assume that agents are risk averse, with constant relative risk aversion parameter 3, over annual incomes. This has only a minor effect. Row 12 shows estimates for larger firing costs ($\kappa = 0.15$). This reduces the benefit of the tenure policy somewhat under the baseline parameter vector, but has a much smaller effect under optimistic parameters. Finally, row 13 uses a lower baseline exit hazard, especially for experienced teachers. Online Appendix Figure 1 shows that many teachers exit the classroom to work in school administration (e.g., as principals). It is not clear that these should be treated as exits for better offers (nor is it clear that they should not be). This change does not much affect the results.

Changes in several parameters at once can be examined by comparing across columns. Column 5 shows that the +0.018 net benefit of the tenure policy turns slightly *negative* under the pessimistic parameters. This largely reflects the reduced exit elasticity ($\zeta = 0.5$, versus $\zeta = 1$ in the baseline). When I adjust the exit elasticity to $\zeta = 1.5$, keeping all other parameters as in the pessimistic scenario (row 8), results are quite similar to those seen in the baseline scenario with $\zeta = 1$ but more favorable values for all other parameters.

IV. Alternative Policies

In this section, I broaden the policy space beyond the simple bonus and tenure contracts considered above. First, I consider combining these contracts with reforms designed to make it easier to dip one's toe in the profession. Second, I present estimates for quantitative changes to the policies, aimed at identifying how the optimal design of each policy varies with the model parameters. Third, I consider alternative retention policies that may make better use of the available information than does a once-and-for-all tenure review.

A. Interactions with Credential Requirements

Employment as a teacher traditionally requires a teaching credential. It is not clear that credentialing programs provide useful training (Boyd et al. 2006; Kane, Rockoff, and Staiger 2008), and the requirement may prevent some potentially able teachers from entering the profession. The barrier would loom largest for those contemplating short teaching careers, so might interact with the tenure contract in particular.

To explore this, I augment the model by requiring prospective teachers to pay a fixed cost, equal to a year's salary under the baseline contract, before entering the profession.²⁵ They demand higher salaries to offset this. I then consider eliminating the entry cost, either alone or in combination with the adoption of a performance-based contract. I adopt the simple but surely incorrect assumption that credentialing programs do not improve teachers' ability or serve a filtering function, so their elimination comes as pure gain.

²⁵There are aspects of teaching other than credential requirements (e.g., the need to invest at the beginning of the career in the development of lesson plans that may be reused later) that can also be interpreted as fixed entry costs. The discussion here applies to these costs as well, but with the important distinction that it is unclear how these costs could be reduced.

TABLE 4—EFFECTS OF CONTRACTS WITH FIXED COSTS OF ENTRY

	No entry cost in baseline (1)	Entry cost in baseline	
		Total impact (2)	Incremental impact of performance contract (3)
Introduce bonus contract	+0.004	+0.004	
Introduce tenure contract	+0.018	+0.018	
Eliminate entry cost			
No other changes		+0.006	
and introduce bonus contract		+0.010	+0.004
and introduce tenure contract		+0.025	+0.018

Notes: All simulations assume a fixed district budget, as in columns 4–5 of Table 2. Simulations in column 2, rows 1–2 assume that entering teachers must pay a fixed entry cost equal to w^0 , both under the baseline contract and under the alternatives. In rows 3–5, teachers pay the entry cost under the baseline contract but not under the alternatives. Table entries represent impacts of the alternative contracts on average teacher productivity, in student-level standard deviations relative to the baseline contract (with a fixed entry cost in column 2, but without one in column 1), incorporating changes in ability, in experience, and in class size.

Results are presented in Table 4. Column 1 repeats results for the baseline case of no entry costs considered above. In column 2, the first rows show the effects of introducing performance contracts when there is a fixed cost to entry. These are identical to the baseline results. The third row shows the effect of eliminating the entry cost under the baseline contract. This allows salaries to be lowered, freeing up enough money to finance 1.7 percent reductions in class size and thereby to improve productivity by 0.006 student-level standard deviations.

Finally, the last rows show the effects of simultaneously eliminating the entry cost and introducing performance contracts. These are equal to three digits to the sum of the separate effects of the two components in isolation. This is not what might have been expected. When compensation is backloaded, as it is with a fixed entry cost and stable growth in post-entry earnings, tenure denials are more costly, and one might expect that teachers—particularly low ability teachers—would demand larger compensation for accepting this risk. Intuition comes from the information structure illustrated in Figure 1. Prospective teachers do not have enough information to forecast their tenure probabilities accurately, so the prospect of a tenure denial weighs nearly equally on high ability and low ability prospective teachers. Both elimination of the entry cost and introduction of the tenure policy thus have effects on entry that are largely uniform across the ability distribution.

There are two important caveats. First, I assume risk neutrality. The policies might interact meaningfully if teachers were risk averse.²⁶ Second, calls for credential reform are often motivated by the idea that many high skilled graduates who foresee highly paid professional careers could be persuaded to teach for a short time if entry costs were low.²⁷ In my model, these graduates have no higher τ than do those with

²⁶ Above, I incorporated risk aversion over annual income. This is unsatisfactory in the presence of a fixed entry cost, which prospective teachers presumably finance out of later earnings. Combining risk aversion with student debt requires a more elaborate model of intertemporal decision making.

²⁷ Teach for America (TFA) is an example: TFA teachers are asked to commit to only two years of teaching, and are assigned classrooms after a five-week training course. Clark et al. (2013) find that TFA teachers are more

worse outside options, so there is little benefit of attracting them. Allowing for a correlation between inside ability and outside options might create an interaction between credential and contract reforms: short-term teachers would not be much affected by tenure decisions, so the salary increases needed to offset the risk of tenure denial for a career teacher would make the profession more attractive to short-term potential teachers. I defer consideration of this to future work.

B. Varying the Bonus and Tenure Contracts

The bonus and tenure policies I have considered thus far are designed to resemble (in both their form and scope) policies that have been implemented by states and school districts. They are far from optimal. In the context of the model here, the optimal pay policy would backload nearly all compensation to the teacher's retirement, when the teacher's ability is known with maximal precision. This is unrealistic for reasons beyond the scope of my stylized model: teachers have consumption needs that make it impossible to wait many years for their salaries, and even if credit were available the risk that one will turn out to have low τ and thus never be paid would not be easily insurable.

First-best policies are thus of little interest. I examine here quantitative variations in the bonus and tenure contracts, while the next subsection examines nontenure retention policies. As above, I assume that the budget is fixed, so that changes in salaries per teacher are offset by changes in the number of classrooms. At the end of this subsection I consider the value of loosening the district budget constraint.

Panel A of Figure 5 shows how the impacts of the bonus contract vary with the size of the annual bonus (expressed as a share of base pay). Not surprisingly, when teachers are risk neutral the effectiveness of the bonus policy increases with the size of the bonus. But the impact of the bonus policy remains small: under baseline parameters, even a bonus equal to 100 percent of base salaries would have a smaller impact than would a policy of denying tenure to 20 percent of second-year teachers. When teachers are risk averse, as with the pessimistic parameters, large bonuses are counterproductive. Here, the maximum impact is achieved with a bonus equal to 21 percent of base pay.

Panel B of Figure 5 varies the threshold for receiving the bonus (expressed as f^B , the share of current teachers who would receive bonuses each year). Regardless of parameter values, the impact of the bonus policy grows with the share of teachers receiving bonuses until this share exceeds 40 percent.

Figure 6 turns to the tenure contract. Here, I vary both the share of teachers denied tenure and the date at which the decision is made; different panels correspond to the different parameter vectors. Panel A of Figure 6 shows that optimal tenure denial rates under the baseline parameters are around 40 percent. The tenure policy is

productive than their more experienced peers. This could indicate a pool of high- τ potential teachers who are unwilling to pay fixed costs to enter. It could also reflect, however, a TFA screening process that successfully selects on τ . Only 20 percent of applicants are selected. In any case, it is not clear that the pool of high ability potential short-term teachers is large enough to scale up dramatically. After years of rapid growth, TFA now produces less than 0.3 percent of all public school teachers. And "no excuses" charter schools (Abdulkadiroğlu et al. 2011) that hire from similar pools serve only a few percent of urban students but are already importantly constrained by labor supply shortages (Wilson 2009).

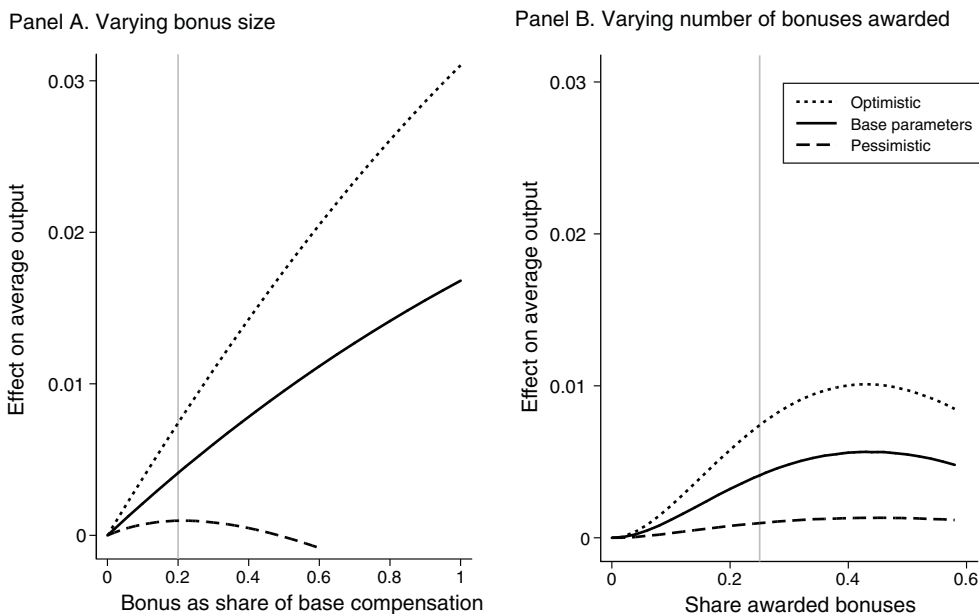


FIGURE 5. VARYING THE PARAMETERS OF THE BONUS CONTRACT

Notes: Panels show the change in average output, relative to the single salary contract and scaled in student-level standard deviations, associated with a shift to the bonus contract with the indicated bonus size (panel A) or threshold for bonus receipt (panel B, with x-axis scaled in terms of the share of current teachers who would receive bonuses each year). Parameters are as indicated in Table 1; base wages are assumed set to fix the total district budget. In panel A, bonuses are awarded to 20 percent of teachers (under the baseline ability distribution); in panel B, the bonus is set to 25 percent of base compensation. Vertical lines indicate contract parameters used for row 1 of Table 3.

notably less effective if tenure decisions are made after only one year, but there is little net benefit (or cost) of waiting more than two years. Longer tenure clocks allow more accurate tenure decisions, but this is offset by the damage done by delaying action on teachers who have already shown themselves to be ineffective. Panels B and C show results for the pessimistic and optimistic parameters. Under the pessimistic parameters (panel B), a 40 percent denial rate is far too high; the optimum is less than 15 percent, and high tenure denial rates are worse than granting tenure to all. Later decisions are preferable here. Under optimistic parameters, by contrast, the optimal policy denies tenure to well over one-half of new teachers, and benefits are not very sensitive to the date of the decision so long as it is made in the second year or later.

The optimal bonus and tenure policies are characterized in panels A and B of Table 5. Tenure policies are uniformly more effective than bonus policies of plausible scale, but the design and impacts of these policies depend critically on the parameters. Even the optimistic parameter values suggest that the benefits of a tenure contract top out around 0.038 student-level standard deviations, and this requires a 30 percent increase in average teacher salaries. The impact on productivity is less than one-half of the impact suggested by Staiger and Rockoff’s (2010) simulation of optimal tenure policies without labor supply responses, in large part due to the class size increases needed to pay for higher salaries. Moreover, where Staiger and

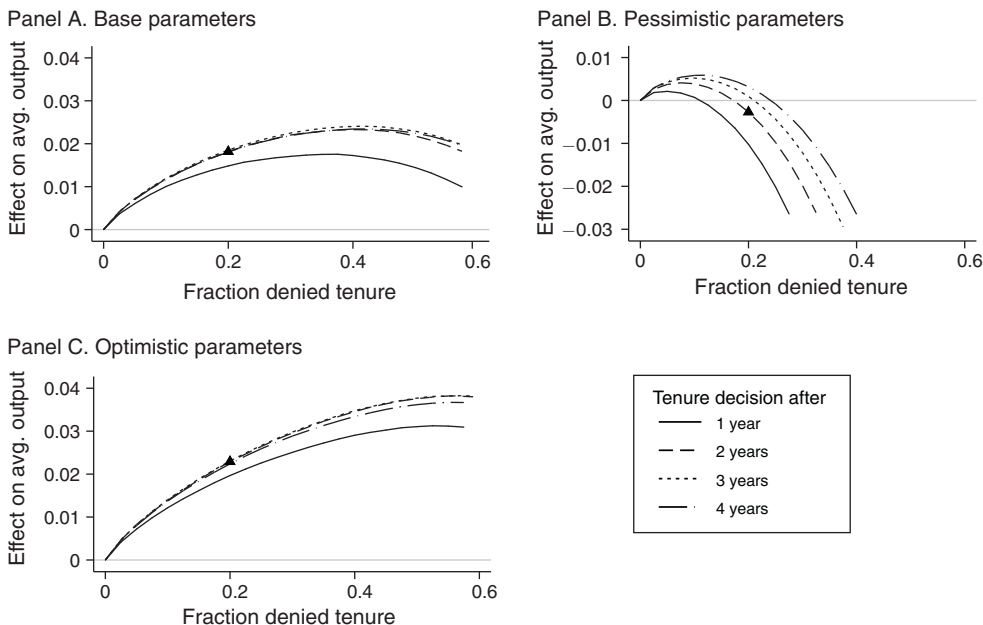


FIGURE 6. VARYING THE TENURE RATE AND TIMING OF TENURE DECISIONS

Notes: Panels show the change in average output, relative to the single salary contract and scaled in student-level standard deviations, associated with a shift to the tenure contract with the indicated tenure denial rate, using different decision dates and parameter values. Parameters are as indicated in Table 1; base wages are assumed set to fix the total district budget. Marked points indicate the contract parameters (20 percent denied tenure, with decisions after the second year) used for row 1 of Table 3.

Rockoff (2010) estimate that tenure decisions should be made after just one year and over 80 percent of new teachers should be dismissed, my results point to later decisions and much higher tenure rates.

The model can also be used to explore interactions between class size and contract terms. Figure A2 in the online Appendix shows estimates for varying tenure rates under both the baseline budget and one that is 5 percent larger.²⁸ Under baseline and optimistic parameters, a 5 percent reduction in class size with zero tenure denials offers nearly identical benefits to a 20 percent tenure denial contract: with baseline parameters, +0.017 for the former and +0.018 for the latter. Moreover, curves for the two budget levels are very nearly parallel as the tenure rate varies, indicating that there is no meaningful interaction between class size reduction and teacher firing. This means that cost-benefit analyses of simultaneous changes in contract terms and district budgets can be conducted by modeling first the impact of a fixed-budget contract change and second the impact of changing class size to balance the new budget, taking each in isolation.²⁹ This is not true under the pessimistic parameters, however. Here, too high a rate of tenure denial does more damage the larger the budget.

²⁸This could increase class size by 5 percent or finance a 15 percent increase in salaries with fixed class size. As Table 2 indicates, the latter would support a 20 percent tenure denial rate.

²⁹For tenure denial rates local to the optimal levels (in Table 5) this follows from the envelope theorem. Figure A2 indicates that effects are additive even when the tenure denial rate is far from the optimum.

TABLE 5—OPTIMAL BONUS AND FIRING POLICIES

	Baseline (1)	Pessimistic (2)	Optimistic (3)
<i>Panel A. Bonuses based on two-year moving average performance</i>			
Varying bonus size, with baseline threshold			
Bonus size	100% [†]	21%	100% [†]
Change in avg. compensation	+1.0%	+0.5%	+2.3%
Change in output	+0.017	+0.001	+0.031
Varying threshold, with baseline size			
Fraction awarded bonus	43%	46%	43%
Change in avg. compensation	+0.004%	+1.0%	+0.1%
Change in output	+0.006	+0.001	+0.010
<i>Panel B. Tenure policies</i>			
Varying date of tenure and fraction denied			
Year of tenure decision	3	5 [†]	3
Fraction denied tenure	41%	13%	57%
Change in avg. compensation	+24%	+10%	+30%
Change in output	+0.024	+0.006	+0.038
Varying fraction denied, fixing decision in year 2			
Fraction denied tenure	40%	7%	56%
Change in avg. compensation	+26%	+8%	+33%
Change in output	+0.023	+0.004	+0.038
<i>Panel C. Ongoing retention decisions</i>			
Based on average performance to date			
Fraction not retained	52%	10%	71%
Change in avg. compensation	+39%	+11%	+53%
Change in output	+0.031	+0.004	+0.054
Based on posterior mean			
Fraction not retained	52%	14%	67%
Change in avg. compensation	+37%	+11%	+45%
Change in output	+0.034	+0.008	+0.055
Based on optimal (static) decision rule			
Fraction not retained	55%	18%	70%
Change in avg. compensation	+36%	+10%	+47%
Change in output	+0.035	+0.008	+0.056

Notes: Daggers represent boundary maxima—I do not consider bonuses larger than 100 percent of base pay or tenure decisions occurring after year 5. All simulations assume a fixed district budget. Changes in average compensation include base salaries, bonuses, and experience premia. Changes in output represent the net change in average productivity, in student-level standard deviations, incorporating changes in ability, experience, and class size. Both changes are relative to the baseline contract.

C. Alternative Firing Policies

Once-and-for-all retention decisions make inefficient use of information: For teachers whose initial performance places them near the retention threshold, error rates would be reduced with a longer probationary period. It is computationally infeasible to solve for the first-best optimal retention rule taking labor supply responses into account. Instead, I consider here three alternative firing contracts that successively better approximate the optimal decision rule in the absence of labor supply responses. The contracts vary in the way that the retention threshold varies with teacher experience.

My first ongoing firing contract conditions retention on the teacher's average performance to date: Any t th-year teacher for whom $\bar{y}_t \equiv \frac{1}{t} \sum_{s=1}^t y_s$ falls below a fixed threshold is fired. I express the threshold in terms of the share of entering teachers

who would be displaced at some point before the end of a 30-year career (assuming that they stay that long, and given the baseline ability distribution). For example, a threshold of -0.26 would lead 20 percent to be displaced at some point before the end of a 30-year career. If the contract were implemented suddenly, 6.8 percent of teachers would be fired immediately, but many of these would have been displaced much earlier had the contract been in place before. Moving forward, 14 percent of new hires would be fired after their first years, 3.1 percent of those who remain would be fired after their second years, and 1.3 percent after their third years. Firing rates would be below 1 percent, but always positive, for $t > 3$.

This contract is in some ways more patient than the tenure contract: 17 percent of fired teachers have more than two years of experience, compared to zero with a two-year tenure clock. But it may not be patient enough. There is more value in retaining a first-year teacher with $y_1 = -0.3$ than in retaining a tenth-year teacher whose average performance to date is so poor, as there is a reasonable prospect that the former was simply unlucky but the latter's performance more likely reflects her true ability. My second ongoing firing contract bases retention decisions on the district's posterior mean of the teacher's ability, $\frac{\sigma_\tau^2}{t^{-1}\sigma_\epsilon^2 + \sigma_\tau^2} \bar{y}_t$. When this contract's threshold is set at a level that would lead 20 percent of teachers to be displaced at some point in their careers, 3.7 percent of teachers (given the current ability distribution) are fired after their first years, 4.4 percent of those who remain are fired after their second years, and 3.1 percent after their third years. Firing probabilities remain above 1 percent through the sixth year.

One might wish to be even more patient than this. The option value of retaining an inexperienced teacher with low posterior mean but high variance is higher than for an experienced teacher with the same posterior mean (so better average performance to date) but low variance—the inexperienced teacher may turn out to be fine, and can always be fired next year if she doesn't. My third contract uses thresholds that vary over time in a way that is optimal from the district's perspective, ignoring labor supply responses.³⁰ This contract displaces only 1.3 percent of first-year teachers, 2.5 percent after their second years, and 2.1 percent after their third years.

Figure A3 in the online Appendix shows the share of teachers at each ability level who are eventually fired under each of the different contracts, when each is calibrated to displace 20 percent of teachers at some point over a 30-year career. The first firing contract does slightly better than the tenure policy at identifying the teachers with the lowest true ability for firing, but the difference is relatively small. The second and third contracts represent more dramatic improvements, firing many more bottom quintile teachers and many fewer teachers outside the bottom quintile. However, this comes at a cost. Figure A4, also in the online Appendix, shows the cumulative firing probability for teachers in the bottom decile of the τ distribution, the second decile, the third and fourth deciles, and the upper six deciles. Although the more patient contracts eventually fire larger shares of the lowest ability teachers,

³⁰The thresholds are computed as the numeric solutions to the district's dynamic optimization problem, assuming that the district pays a firing cost that is proportional to the number of years of labor supply foregone and ignores labor supply responses. Note that I fix the overall firing rate by setting the firing cost; if it is set above the district's true shadow cost, the firing rate is suboptimally low.

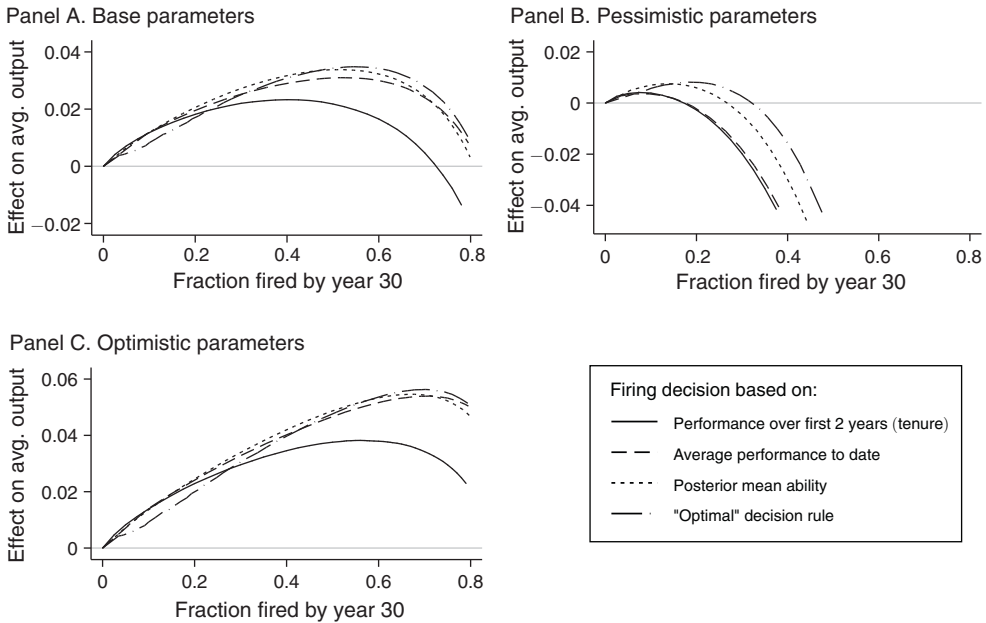


FIGURE 7. COMPARING UP-OR-OUT TENURE TO ONGOING RETENTION DECISIONS

Notes: Panels show the change in average output, relative to the single salary contract and scaled in student-level standard deviations, associated with shifts to contracts incorporating ongoing retention decisions. The decision rules are described in the text. x -axis measures the fraction of teachers who would be fired before the end of a 30-year career, given the ability distribution of current entering teachers. Parameters are as indicated in Table 1; base wages are assumed set to fix the total district budget.

they wait longer to do so—substantially so for second decile teachers under the “optimal” decision rule.

Figure 7 presents the impacts of the alternative decision rules at different scales. It shows that the ongoing firing contracts support higher firing rates but that at modest firing rates they have similar impacts to the tenure policy. As expected, patience is most useful when the firing rate is set to maximize output, but when the firing rate is kept below its optimal level it can be better to make faster, more error prone decisions than to wait to optimally distinguish among teachers just above and below the desired threshold.

Results for the optimal scale of each contract type are presented in the lower panel of Table 5. All but one of the ongoing retention contracts outperforms the tenure contract; in each case, more patience allows larger net productivity improvements than do the less patient contracts, usually with higher firing rates but lower salaries. However, the optimal decision rules and policy impacts are quite sensitive to the model parameters. Where under the pessimistic parameters the firing rate never exceeds 18 percent, the baseline parameters yield optimal firing rates as high as 55 percent, and the optimistic parameters yield firing thresholds that would displace as many as 71 percent of teachers before the end of a 30-year career. The alternative contracts would yield net productivity improvements ranging from just over 2 percent of a teacher-level standard deviation, for the least patient contract under pessimistic parameters, to nearly 38 percent, for the most patient decision rule under optimistic parameters.

V. Discussion

The simulations presented here suggest that the effects of policies aimed at improving teacher productivity will depend importantly on their interactions with the teacher labor market. If prospective teachers are uncertain about their own abilities or if their labor supply is less than perfectly elastic, performance-based compensation and retention policies each require substantial increases in teacher salaries. These matter to the evaluation of alternative contracts. Financing them with a fixed budget requires class size increases that offset about one-half of the gross benefits of the alternative contracts.

Despite the high costs, both bonus and firing policies can be cost effective. Indeed, recognition of the labor market effects can make these policies even more effective than when these effects are ignored, as the accompanying salary increases help to attract and retain high ability teachers. Policy design is important, however, as cost-effectiveness varies substantially with the specifics of the contract. I find, for example, that when firing rates are very high the option to fire experienced teachers has substantial value, but when firing rates are lower, early, irrevocable tenure decisions are approximately optimal.

The gains from improved policies could be substantial. Under my baseline parameters, a fixed-budget increase in the tenure denial rate from 0 to 20 percent would raise output by 0.12 teacher-level standard deviations (column 4 of Table 2). Chetty, Friedman, and Rockoff's (2014b) estimates of the association between teacher value-added and students' later earnings, discussed in Section IIIB, suggest that this would yield present-value benefits of about \$24,000 per teacher per year.³¹ Even if the true effects are a fraction of this (as suggested by Rothstein 2014), a good deal is at stake.

There are several important caveats, however. First, the results depend importantly on parameter values. Policies that are optimal under one set of parameter values can be harmful under other plausible parameters. Under the pessimistic parameter vector, an increase in the tenure denial rate from 0 to 20 percent would *reduce* output. Results are particularly sensitive to the labor supply elasticity and the degree of foreknowledge that prospective teachers possess, and future research should aim to uncover these parameters.

Second, the analysis relies on a best case view of the potential for teacher performance assessment. I assume that performance measures are unbiased, cover the full range of desired outputs, and are not subject to "influence activities" that raise measured performance without raising true productivity. None of these is very plausible.

Consider first the case where output is multidimensional and the performance measure captures only one of the dimensions. For example, the performance measure might focus on cognitive skills though teachers also teach noncognitive skills,

³¹ Chetty, Friedman, and Rockoff's (2014b) results also imply that "traditional" policy changes would have large impacts. For example, one could obtain similar benefits by increasing the school budget by 5 percent (or \$7,500 per teacher) under a zero-firing contract, implying a benefit-cost ratio around 3. See also Chetty et al. (2011) and the discussion in Section IIIB, above. The results in Section IVB imply that the effects of contract and budget changes are additive, so financing the tenure contract with increased spending rather than with class size increases would have effects nearly twice as large as those in the text. My model can also be used to estimate the optimal spending level (though this requires extrapolating the class size effect far beyond the STAR experiment). Under my base parameters, the optimal budget would be more than triple the status quo.

or it might focus only on certain subjects, or weight test-taking skills too heavily. The policies I consider here will improve teacher ability on unmeasured dimensions in proportion to the correlation of ability in that dimension with that in the measured dimension. Value-added measures are correlated only about 0.4 (after adjustment for attenuation due to sampling error) across different tests in the same subject (see, e.g., Bill and Melinda Gates Foundation 2010; Rothstein 2011). Correlations between value-added scores and other performance measures (e.g., classroom observations) are even lower (Bill and Melinda Gates Foundation 2012; Rothstein and Mathis 2013).

Matters are even worse if teachers can affect their measured performance, either by reallocating effort between measured and unmeasured outputs (sometimes known as “goal distortion”) or by manipulating the performance measurement process. High-stakes evaluations can be counterproductive in this case (Baker 1992, 2002; Holmström and Milgrom 1991). There is evidence that teachers can improve their measured value-added by reducing the attention paid to nontested topics and subjects, teaching to the test, arranging to have the right students, or outright cheating, and that teachers faced with high-stakes incentives will respond at least in part in these ways (e.g., Campbell 1979; Neal and Schanzenbach 2010; Rothstein 2010; Carrell and West 2010). Rothstein (2012) finds that the benefits of performance-based contracts are quite sensitive to the potential for goal distortion and manipulation. A high priority topic for future research must be the degree to which productivity measures become corrupted when the stakes are raised (Rothstein 2011).

Finally, there are many aspects of the teaching profession omitted from my stylized model. I do not account for the possibility that teachers may be self-selected for unusual risk aversion; for the social status of teachers relative to other professions; or for the potential for high-stakes evaluations to undermine cooperation among teachers and principals. Moreover, I assume that new teachers recruited under alternative contracts would come from the same general population as do current teachers and do not allow for the possibility, sometimes raised in discussions of teacher quality, that there exists a separate pool of high ability potential teachers who would not consider teaching under current conditions. These issues are not well enough understood to incorporate into my quantitative model; extensions of the model to allow for them are left as a subject for future research.

These caveats aside, the analysis here demonstrates that clear thinking about the impact of teacher quality policy changes requires a model of the roles of imperfect information, teacher salaries, and labor supply decisions. Even in my best case scenarios, alternative teacher contracts have more modest impacts on student achievement than has often been promised. None of the contracts considered here would raise average productivity by more than 40 percent of a standard deviation. More plausible parameters and policies yield improvements that are generally less than half that size. These kinds of benefits would be most welcome, but would not represent fundamental changes in our education system.

REFERENCES

- Abdulkadiroğlu, Atila, Joshua D. Angrist, Susan M. Dynarski, Thomas J. Kane, and Parag A. Pathak.** 2011. "Accountability and Flexibility in Public Schools: Evidence from Boston's Charters and Pilots." *Quarterly Journal of Economics* 126 (2): 699–748.
- Adda, Jerome, Christian Dustmann, Costas Meghir, and Jean-Marc Robin.** 2013. "Career Progression, Economic Downturns, and Skills." National Bureau of Economic Research Working Paper 18832.
- Baker, George P.** 1992. "Incentive Contracts and Performance Measurement." *Journal of Political Economy* 100 (3): 598–614.
- Baker, George.** 2002. "Distortion and Risk in Optimal Incentive Contracts." *Journal of Human Resources* 37 (4): 728–51.
- Ballou, Dale.** 1996. "Do Public Schools Hire the Best Applicants?" *Quarterly Journal of Economics* 111 (1): 97–133.
- Bill and Melinda Gates Foundation.** 2010. *Learning About Teaching: Initial Findings from the Measures of Effective Teaching Project*. Seattle: MET Project.
- Bill and Melinda Gates Foundation.** 2012. *Gathering Feedback for Teaching: Combining High-Quality Observations with Student Surveys and Achievement Gains*. Seattle: MET Project.
- Boyd, Donald, Pamela Grossman, Hamilton Lankford, Susanna Loeb, and James Wyckoff.** 2006. "How Changes in Entry Requirements Alter the Teacher Workforce and Affect Student Achievement." *Education Finance and Policy* 1 (2): 176–216.
- Boyd, Donald, Hamilton Lankford, Susanna Loeb, and James Wyckoff.** 2011. "Teacher Layoffs: An Empirical Illustration of Seniority versus Measures of Effectiveness." *Education Finance and Policy* 6 (3): 439–54.
- Campbell, Donald T.** 1979. "Assessing the Impact of Planned Social Change." *Evaluation and Program Planning* 2 (1): 67–90.
- Carrell, Scott E., and James E. West.** 2010. "Does Professor Quality Matter? Evidence from Random Assignment of Students to Professors." *Journal of Political Economy* 118 (3): 409–32.
- Chandler, Michael Alison.** 2007. "Schools Pinched in Hiring." *The Washington Post*. June 24.
- Chetty, Raj, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan.** 2011. "How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project STAR." *Quarterly Journal of Economics* 126 (4): 1593–660.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff.** 2014a. "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates." *American Economic Review* 104 (9): 2593–632.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff.** 2014b. "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood." *American Economic Review* 104 (9): 2633–79.
- Chingos, Matthew M., and Martin R. West.** 2012. "Do More Effective Teachers Earn More Outside the Classroom?" *Education Finance and Policy* 7 (1): 8–43.
- Clark, Melissa A., Hanley S. Chiang, Tim Silva, Sheena McConnell, Kathy Sonnenfeld, Anastasia Erbe, and Michael Puma.** 2013. *The Effectiveness of Secondary Math Teachers from Teach For America and the Teaching Fellows Programs*. Washington, DC: National Center for Education Evaluation, US Department of Education.
- Clotfelter, Charles, Elizabeth Glennie, Helen Ladd, and Jacob Vigdor.** 2008. "Would Higher Salaries Keep Teachers in High-Poverty Schools? Evidence from a Policy Intervention in North Carolina." *Journal of Public Economics* 92 (5–6): 1352–70.
- Clotfelter, Charles T., Helen F. Ladd, and Jacob L. Vigdor.** 2007. "How and Why do Teacher Credentials Matter for Student Achievement?" National Bureau of Economic Research Working Paper 12828.
- Corcoran, Sean P., William N. Evans, and Robert M. Schwab.** 2004. "Changing Labor-Market Opportunities for Women and the Quality of Teachers, 1957–2000." *American Economic Review* 94 (2): 230–5.
- Dee, Thomas, and James Wyckoff.** 2013. "Incentives, Selection, and Teacher Performance: Evidence from IMPACT." National Bureau of Economic Research Working Paper 19529.
- Dolton, Peter, and Wilbert van der Klaauw.** 1999. "The Turnover of Teachers: A Competing Risks Explanation." *Review of Economics and Statistics* 81 (3): 543–50.
- Duflo, Esther, Rema Hanna, and Stephen P. Ryan.** 2012. "Incentives Work: Getting Teachers to Come to School." *American Economic Review* 102 (4): 1241–78.

- Figlio, David N.** 2002. "Can Public Schools Buy Better-Qualified Teachers?" *Industrial and Labor Relations Review* 55 (4): 686–99.
- Flyer, Fredrick, and Sherwin Rosen.** 1997. "The New Economics of Teachers and Education." *Journal of Labor Economics* 15 (1): S104–39.
- Fryer, Roland G.** 2013. "Teacher Incentives and Student Achievement: Evidence from New York City Public Schools." *Journal of Labor Economics* 31 (2): 373–407.
- Fryer, Roland G., Jr., Steven D. Levitt, John List, and Sally Sadoff.** 2012. "Enhancing the Efficacy of Teacher Incentives through Loss Aversion: A Field Experiment." National Bureau of Economic Research Working Paper 18237.
- Gates, Bill.** 2011. "How Teacher Development Could Revolutionize Our Schools." *The Washington Post*. February 28.
- Glazerman, Steven, Hanley Chiang, Alison Wellington, Jill Constantine, and Dan Player.** 2011. *Impacts of Performance Pay Under the Teacher Incentive Fund: Study Design Report*. Washington, DC: Mathematica Policy Research Report.
- Glewwe, Paul, Nauman Ilias, and Michael Kremer.** 2010. "Teacher Incentives." *American Economic Journal: Applied Economics* 2 (3): 205–27.
- Goldhaber, Dan, Betheny Gross, and Daniel Player.** 2011. "Teacher Career Paths, Teacher Quality, and Persistence in the Classroom: Are Public Schools Keeping Their Best?" *Journal of Policy Analysis and Management* 30 (1): 57–87.
- Goodman, Sarena F., and Lesley J. Turner.** 2013. "The Design of Teacher Incentive Pay and Educational Outcomes: Evidence from the New York City Bonus Program." *Journal of Labor Economics* 31 (2): 409–20.
- Gordon, Robert, Thomas J. Kane, and Douglas O. Staiger.** 2006. *Identifying Effective Teachers Using Performance on the Job*. Washington, DC: Brookings Institution Press.
- Hanushek, Eric A., and Steven G. Rivkin.** 2006. "Teacher Quality." In *Handbook of the Economics of Education*. Vol. 2, edited by Eric Hanushek and Finis Welch, 1051–78. Amsterdam: Elsevier B.V.
- Hess, Rick.** 2010. "Missing the POINT: Tomorrow's Big Merit Pay Study Will Tell Us...Nothing." *Education Week's Blog: Rick Hess Straight Up*. September 20. http://blogs.edweek.org/edweek/rick_hess_straight_up/2010/09/missing_the_point_tomorrows_big_merit_pay_study_will_tell_us_nothing.html (accessed November 1, 2014).
- Hiatt, Fred.** 2009. "How Bill Gates Would Repair Our Schools." *The Washington Post*. March 30.
- Holmström, Bengt, and Paul Milgrom.** 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, and Organization* 7: 24–52.
- Kane, Thomas J., Daniel F. McCaffrey, Trey Miller, and Douglas O. Staiger.** 2013. *Have We Identified Effective Teachers? Validating Measures of Effective Teaching Using Random Assignment*. Seattle: MET Project.
- Kane, Thomas J., Jonah E. Rockoff, and Douglas O. Staiger.** 2008. "What Does Certification Tell Us about Teacher Effectiveness? Evidence from New York City." *Economics of Education Review* 27 (6): 615–31.
- Keane, Michael P., and Kenneth I. Wolpin.** 1997. "The Career Decisions of Young Men." *Journal of Political Economy* 105 (3): 473–522.
- Klein, Joel, Michelle Rhee, Peter C. Gorman, Ron Huberman, Carol R. Johnson, Andrés A. Alonso, Tom Boasberg et al.** 2010. "How to Fix Our Schools: A Manifesto by Joel Klein, Michelle Rhee and Other Education Leaders." *The Washington Post*. October 10.
- Krieg, John M.** 2006. "Teacher Quality and Attrition." *Economics of Education Review* 25 (1): 13–27.
- Krueger, Alan B.** 1999. "Experimental Estimates of Education Production Functions." *Quarterly Journal of Economics* 114 (2): 497–532.
- Lavy, Victor.** 2002. "Evaluating the Effect of Teachers' Group Performance Incentives on Pupil Achievement." *Journal of Political Economy* 110 (6): 1286–317.
- Lazear, Edward P.** 2000. "Performance Pay and Productivity." *American Economic Review* 90 (5): 1346–61.
- Lazear, Edward P.** 2003. "Teacher Incentives." *Swedish Economic Policy Review* 10 (2): 179–214.
- Manning, Alan.** 2005. *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton: Princeton University Press.
- Muralidharan, Karthik, and Venkatesh Sundararaman.** 2011. "Teacher Performance Pay: Experimental Evidence from India." *Journal of Political Economy* 119 (1): 39–77.
- Murnane, Richard J., and David K. Cohen.** 1986. "Merit Pay and the Evaluation Problem: Why Most Merit Pay Plans Fail and a Few Survive." *Harvard Educational Review* 56 (1): 1–17.
- Murnane, Richard J., and Randall J. Olsen.** 1989. "The Effect of Salaries and Opportunity Costs on Duration in Teaching: Evidence from Michigan." *Review of Economics and Statistics* 71 (2): 347–52.

- Neal, Derek, and Diane Whitmore Schanzenbach. 2010. "Left Behind by Design: Proficiency Counts and Test-Based Accountability." *Review of Economics and Statistics* 92 (2): 263–83.
- Prendergast, Canice. 1999. "The Provision of Incentives in Firms." *Journal of Economic Literature* 37 (1): 7–63.
- Ransom, Michael R., and David P. Sims. 2010. "Estimating the Firm's Labor Supply Curve in a 'New Monopsony' Framework: Schoolteachers in Missouri." *Journal of Labor Economics* 28 (2): 331–55.
- Rivkin, Steven G., Eric A. Hanushek, and John F. Kain. 2005. "Teachers, Schools, and Academic Achievement." *Econometrica* 73 (2): 417–58.
- Rockoff, Jonah E., Brian A. Jacob, Thomas J. Kane, and Douglas O. Staiger. 2011. "Can You Recognize an Effective Teacher When You Recruit One?" *Education Finance and Policy* 6 (1): 43–74.
- Rothstein, Jesse. 2010. "Teacher Quality in Educational Production: Tracking, Decay, and Student Achievement." *Quarterly Journal of Economics* 125 (1): 175–214.
- Rothstein, Jesse. 2011. *Review of 'Learning About Teaching: Initial Findings from the Measures Of Effective Teaching Project'*. Boulder, CO: National Education Policy Center.
- Rothstein, Jesse. 2012. "Teacher Quality Policy When Supply Matters." National Bureau of Economic Research Working Paper 18419.
- Rothstein, Jesse. 2014. "Revisiting the Impacts of Teachers." Unpublished.
- Rothstein, Jesse. 2015. "Teacher Quality Policy When Supply Matters: Dataset." *American Economic Review*. <http://dx.doi.org/10.1257/aer.20121242>.
- Rothstein, Jesse, and William J. Mathis. 2013. *Review of Two Culminating Reports from the MET Project*. Boulder, CO: National Education Policy Center.
- Sass, Tim R. 2008. *The Stability of Value-Added Measures of Teacher Quality and Implications for Teacher Compensation Policy*. Washington, DC: National Center for Analysis of Longitudinal Data in Education Research.
- Springer, Matthew G., Dale Ballou, Laura Hamilton, Vi-Nhuan Le, J. R. Lockwood, Daniel F. McCaffrey, Matthew Pepper, and Brian M. Stecher. 2010. "Teacher Pay for Performance: Experimental Evidence from the Project On Incentives in Teaching." Nashville: National Center on Performance Incentives.
- Staiger, Douglas O., and Jonah E. Rockoff. 2010. "Searching for Effective Teachers with Imperfect Information." *Journal of Economic Perspectives* 24 (3): 97–118.
- Taylor, Eric S., and John H. Tyler. 2012. "The Effect of Evaluation on Teacher Performance." *American Economic Review* 102 (7): 3628–51.
- Tincani, Michela. 2011. "Teacher Labor Markets, School Vouchers and Student Cognitive Achievement: Evidence from Chile." Unpublished.
- von Wachter, Till, Jae Song, and Joyce Manchester. 2009. "Long-Term Earnings Losses Due to Mass Layoffs During the 1982 Recession: An Analysis Using U.S. Administrative Data from 1974 to 2004." Unpublished.
- West, Martin, and Matthew Chingos. 2009. "Teacher Effectiveness, Mobility, and Attrition in Florida." In *Performance Incentives: Their Growing Impact on American K12 Education*, edited by Matthew G. Springer, 251–72. Washington, DC: Brookings Institution Press.
- Wilson, Steven F. 2009. "Success at Scale in Charter Schooling." American Enterprise Institute Education Outlook 3.
- Winters, Marcus A., and Joshua M. Cowen. 2013a. "Who Would Stay, Who Would Be Dismissed? An Empirical Consideration of Value-Added Teacher Retention Policies." *Educational Researcher* 42 (6): 330–7.
- Winters, Marcus A., and Joshua M. Cowen. 2013b. "Would a Value-Added System of Retention Improve the Distribution of Teacher Quality? A Simulation of Alternative Policies." *Journal of Policy Analysis and Management* 32 (3): 634–54.
- Wiswall, Matthew. 2007. "Licensing and Occupational Sorting in the Market for Teachers." Unpublished.

This article has been cited by:

1. Il Hwan Chung, Jongmin Shon, Moonyoung Eom. 2024. Collective bargaining agreements and teacher sorting: evidence from New York. *Asia Pacific Education Review* 25:1, 5-18. [[Crossref](#)]
2. Kevin Ng. 2023. Identifying high-quality teachers. *Education Economics* 2, 1-28. [[Crossref](#)]
3. Frank Perrone, Coby V. Meyers. 2023. A Review of the Research on Teacher Hiring in the United States (2001–2021). *Journal of Education Human Resources* 111. . [[Crossref](#)]
4. Beng Huat See, Stephen Gorard, Rebecca Morris. Tackling teacher shortages: doing the right thing at the right time and with the right evidence 655-670. [[Crossref](#)]
5. Joshua Bleiberg. Revisiting teacher evaluation a decade after reforms 54-61. [[Crossref](#)]
6. Eric A. Hanushek, Jin Luo, Andrew Morgan, Minh Nguyen, Ben Ost, Steven Rivkin, Ayman Shakeel. 2023. The Effects of Comprehensive Educator Evaluation and Pay Reform on Achievement. *SSRN Electronic Journal* 105. . [[Crossref](#)]
7. Michael Gilraine, James Graham, Angela Zheng. 2023. Public Education and Intergenerational Housing Wealth Effects. *SSRN Electronic Journal* 107. . [[Crossref](#)]
8. Jo Blanden, Matthias Doepke, Jan Stuhler. Educational inequality 405-497. [[Crossref](#)]
9. Andrew Bacher-Hicks, Cory Koedel. Estimation and interpretation of teacher value added in research applications 93-134. [[Crossref](#)]
10. Eric S. Taylor. Teacher evaluation and training 61-141. [[Crossref](#)]
11. Danielle Victoria Handel, Eric A. Hanushek. US school finance: Resources and outcomes 143-226. [[Crossref](#)]
12. Danuta Piróg. 2022. Powstawanie niedoborów kadrowych w szkolnictwie. Przegląd stanu badań oraz rekomendacje dla lokalnej polityki w zakresie usług edukacyjnych. *Rozwój Regionalny i Polityka Regionalna* 15:61. . [[Crossref](#)]
13. Kaitlin P. Anderson, Joshua M. Cowen, Katharine O. Strunk. 2022. The Impact of Teacher Labor Market Reforms on Student Achievement: Evidence from Michigan. *Education Finance and Policy* 17:4, 666-692. [[Crossref](#)]
14. David D. Liebowitz, Lorna Porter, Dylan Bragg. 2022. The Effects of Higher-Stakes Teacher Evaluation on Office Disciplinary Referrals. *Journal of Research on Educational Effectiveness* 15:3, 475-509. [[Crossref](#)]
15. James Cowan, Dan Goldhaber, Roddy Theobald. 2022. Performance Evaluations as a Measure of Teacher Effectiveness When Implementation Differs: Accounting for Variation across Classrooms, Schools, and Districts. *Journal of Research on Educational Effectiveness* 15:3, 510-531. [[Crossref](#)]
16. Nathan Barrett, Deven Carlson, Douglas N. Harris, Jane Arnold Lincove. 2022. When the Walls Come Down: Evidence on Charter Schools' Ability to Keep Their Best Teachers Without Unions and Certification Rules. *Educational Evaluation and Policy Analysis* 44:2, 283-312. [[Crossref](#)]
17. Michael Bates, Michael Dinerstein, Andrew C. Johnston, Isaac Sorkin. 2022. Teacher Labor Market Equilibrium and Student Achievement. *SSRN Electronic Journal* 107. . [[Crossref](#)]
18. Andrew Johnston. 2022. Preferences, Selection, and the Structure of Teacher Pay. *SSRN Electronic Journal* 92. . [[Crossref](#)]
19. Michael Bates, Michael Dinerstein, Andrew Johnston, Isaac Sorkin. 2022. Teacher Labor Market Equilibrium and Student Achievement. *SSRN Electronic Journal* 107. . [[Crossref](#)]
20. Jo Blanden, Matthias Doepke, Jan Stuhler. 2022. Educational Inequality. *SSRN Electronic Journal* 127. . [[Crossref](#)]

21. Mauricio G. Villena, Juan Díaz, Rafael Sanchez, Gabriel Villarroel. 2022. Estimating the Impact of Teachers' Quality on Tertiary Education Attendance: Value Added Versus Evaluation Tests. *SSRN Electronic Journal* 3. . [[Crossref](#)]
22. Stephen Gibbons, Vincenzo Scrutinio, Shqiponja Telhaj. 2021. Teacher turnover: Effects, mechanisms and organisational responses. *Labour Economics* 73, 102079. [[Crossref](#)]
23. Nathan Barrett, Katharine O. Strunk, Jane Lincove. 2021. When tenure ends: the short-run effects of the elimination of Louisiana's teacher employment protections on teacher exit and retirement. *Education Economics* 29:6, 559-579. [[Crossref](#)]
24. Simon Burgess, Shenila Rawal, Eric S. Taylor. 2021. Teacher Peer Observation and Student Test Scores: Evidence from a Field Experiment in English Secondary Schools. *Journal of Labor Economics* 39:4, 1155-1186. [[Crossref](#)]
25. Barbara Biasi. 2021. The Labor Market for Teachers under Different Pay Schemes. *American Economic Journal: Economic Policy* 13:3, 63-102. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
26. David D. Liebowitz. 2021. Teacher evaluation for accountability and growth: Should policy treat them as complements or substitutes?. *Labour Economics* 71, 102024. [[Crossref](#)]
27. Clare Leaver, Owen Ozier, Pieter Serneels, Andrew Zeitlin. 2021. Recruitment, Effort, and Retention Effects of Performance Contracts for Civil Servants: Experimental Evidence from Rwandan Primary Schools. *American Economic Review* 111:7, 2213-2246. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
28. Sebastian Bauhoff, Eeshani Kandpal. Information, Loss Framing, and Spillovers in Pay-for-Performance Contracts 53, . [[Crossref](#)]
29. Aaron L. Schwartz. 2021. Accuracy versus Incentives. *American Journal of Health Economics* 7:3, 333-360. [[Crossref](#)]
30. Richard Murphy, Felix Weinhardt, Gill Wyness. 2021. Who teaches the teachers? A RCT of peer-to-peer observation and feedback in 181 schools. *Economics of Education Review* 82, 102091. [[Crossref](#)]
31. Thomas S. Dee, Jessalynn James, Jim Wyckoff. 2021. Is Effective Teacher Evaluation Sustainable? Evidence from District of Columbia Public Schools. *Education Finance and Policy* 16:2, 313-346. [[Crossref](#)]
32. Jeanne M. Powers, Kathryn P. Chapman. 2021. Poor Kids versus Bad Teachers: Vergara v. California and the Social Construction of Teachers. *Teachers College Record: The Voice of Scholarship in Education* 123:4, 1-26. [[Crossref](#)]
33. José Felipe Martínez, María Paz Fernández. Teacher Evaluation with Multiple Indicators: Conceptual and Methodological Considerations Regarding Validity 373-394. [[Crossref](#)]
34. Julie Berry Cullen, Cory Koedel, Eric Parsons. 2021. The Compositional Effect of Rigorous Teacher Evaluation on Workforce Quality. *Education Finance and Policy* 16:1, 7-41. [[Crossref](#)]
35. Andrew Zeitlin. 2020. Teacher Turnover in Rwanda. *Journal of African Economies* 34. . [[Crossref](#)]
36. Li Han, Jiaxin Xie. 2020. Can conditional grants attract better students? Evidence from Chinese teachers' colleges. *Economics of Education Review* 78, 102034. [[Crossref](#)]
37. Matthew A. Kraft, Eric J. Brunner, Shaun M. Dougherty, David J. Schwegman. 2020. Teacher accountability reforms and the supply and quality of new teachers. *Journal of Public Economics* 188, 104212. [[Crossref](#)]
38. Samuel Berlinski, Alejandra Ramos. 2020. Teacher mobility and merit pay: Evidence from a voluntary public award program. *Journal of Public Economics* 186, 104186. [[Crossref](#)]
39. Nava Ashraf, Oriana Bandiera, Edward Davenport, Scott S. Lee. 2020. Losing Prosociality in the Quest for Talent? Sorting, Selection, and Productivity in the Delivery of Public Services. *American Economic Review* 110:5, 1355-1394. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]

40. Dayna Jean DeFeo, Matthew Berman, Diane Hirshberg. 2020. Statute and Implementation: How Phantom Policies Affect Tenure Value and Support. *Educational Policy* 34:2, 350-376. [[Crossref](#)]
41. Christine M. Dabbs. 2020. Restricting seniority as a factor in public school district layoffs: Analyzing the impact of state legislation on graduation rates. *Economics of Education Review* 74, 101926. [[Crossref](#)]
42. Ajree D. Malawani, Salahudin Salahudin, Zuly Qodir, Mohammad Jafar Loilatu, Achmad Nurmandi. The Evolution of “GOJEK” as an Indonesian Urban Mobile Ride Hailing Model Study Case: Public and Government Regulatory Responses on Urban Mobile Ride Hailing 429-438. [[Crossref](#)]
43. Jessalynn James, James Wyckoff. Teacher labor markets: An overview 355-370. [[Crossref](#)]
44. Andrew C. Johnston. 2020. Teacher Preferences, Working Conditions, and Compensation Structure. *SSRN Electronic Journal* 92. . [[Crossref](#)]
45. Eric Brunner, Joshua M. Cowen, Katharine O. Strunk, Steven Drake. 2019. Teacher Labor Market Responses to Statewide Reform: Evidence From Michigan. *Educational Evaluation and Policy Analysis* 41:4, 403-425. [[Crossref](#)]
46. Steven Drake, Amy Auletto, Joshua M. Cowen. 2019. Grading Teachers: Race and Gender Differences in Low Evaluation Ratings and Teacher Employment Outcomes. *American Educational Research Journal* 56:5, 1800-1833. [[Crossref](#)]
47. Daniel Gibbs. 2019. Selection rates and bureaucratic performance. *Economics of Governance* 20:2, 159-181. [[Crossref](#)]
48. Gregory Camilli. Teacher Influences and Program Effectiveness in Early Childhood Education 52-73. [[Crossref](#)]
49. E. Jason Baron. 2019. Union Reform, Performance Pay, and New Teacher Supply: Evidence from Wisconsin's Act 10. *SSRN Electronic Journal* . [[Crossref](#)]
50. Nirav Mehta. 2019. Measuring quality for use in incentive schemes: The case of “shrinkage” estimators. *Quantitative Economics* 10:4, 1537-1577. [[Crossref](#)]
51. Pablo Munoz, Mounu Prem. 2019. Productivity and Sectoral Allocation: The Labor Market of School Principals. *SSRN Electronic Journal* 25. . [[Crossref](#)]
52. Lindsay Bell Weixler, Douglas N. Harris, Nathan Barrett. 2018. Teachers’ Perspectives on the Learning and Work Environments Under the New Orleans School Reforms. *Educational Researcher* 47:8, 502-515. [[Crossref](#)]
53. Kendall Deas. High-performing Teachers, Student Achievement, and Equity as an Outcome of Educational Supervision 555-574. [[Crossref](#)]
54. Brian A. Jacob, Jonah E. Rockoff, Eric S. Taylor, Benjamin Lindy, Rachel Rosen. 2018. Teacher applicant hiring and teacher performance: Evidence from DC public schools. *Journal of Public Economics* 166, 81-97. [[Crossref](#)]
55. Kirabo Jackson, Alexey Makarin. 2018. Can Online Off-the-Shelf Lessons Improve Student Outcomes? Evidence from a Field Experiment. *American Economic Journal: Economic Policy* 10:3, 226-254. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
56. Ben Backes, James Cowan, Dan Goldhaber, Cory Koedel, Luke C. Miller, Zeyu Xu. 2018. The common core conundrum: To what extent should we worry that changes to assessments will affect test-based measures of teacher performance?. *Economics of Education Review* 62, 48-65. [[Crossref](#)]
57. Richard Murphy, Felix Weinhardt, Gill Wyness. 2018. Who Teaches the Teacher? A RCT of Peer-to-Peer Observation and Feedback in 181 Schools. *SSRN Electronic Journal* 25. . [[Crossref](#)]
58. Jesse Rothstein. 2017. Measuring the Impacts of Teachers: Comment. *American Economic Review* 107:6, 1656-1684. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]

59. Raj Chetty, John N. Friedman, Jonah E. Rockoff. 2017. Measuring the Impacts of Teachers: Reply. *American Economic Review* **107**:6, 1685-1717. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
60. Cory Koedel, Jiayi Li, Matthew G. Springer, Li Tan. 2017. The Impact of Performance Ratings on Job Satisfaction for Public School Teachers. *American Educational Research Journal* **54**:2, 241-278. [[Crossref](#)]
61. Dhushyanth Raju. 2017. Public School Teacher Management in Sri Lanka. *South Asia Economic Journal* **18**:1, 39-63. [[Crossref](#)]
62. Melinda Adnot, Thomas Dee, Veronica Katz, James Wyckoff. 2017. Teacher Turnover, Teacher Quality, and Student Achievement in DCPS. *Educational Evaluation and Policy Analysis* **39**:1, 54-76. [[Crossref](#)]
63. Dan Goldhaber, Richard Startz. 2017. On the Distribution of Worker Productivity: The Case of Teacher Effectiveness and Student Achievement. *Statistics and Public Policy* **4**:1, 1-12. [[Crossref](#)]
64. Barbara Biasi. 2017. Unions, Salaries, and the Market for Teachers: Evidence from Wisconsin. *SSRN Electronic Journal* . [[Crossref](#)]
65. Patrick Kline, Christopher R. Walters. 2016. Evaluating Public Programs with Close Substitutes: The Case of Head Start*. *The Quarterly Journal of Economics* **131**:4, 1795-1848. [[Crossref](#)]
66. Aaron Chalfin, Oren Danieli, Andrew Hillis, Zubin Jelveh, Michael Luca, Jens Ludwig, Sendhil Mullainathan. 2016. Productivity and Selection of Human Capital with Machine Learning. *American Economic Review* **106**:5, 124-127. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
67. Raj Chetty, John N. Friedman, Jonah Rockoff. 2016. Using Lagged Outcomes to Evaluate Bias in Value-Added Models. *American Economic Review* **106**:5, 393-399. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
68. Dan Goldhaber, Joe Walch. 2016. Teacher tenure. *Phi Delta Kappan* **97**:6, 8-15. [[Crossref](#)]
69. Sean Corcoran. Section I Discussion: How Do Educators Use Student Growth Measures in Practice? 153-165. [[Crossref](#)]
70. Cristina Giorgantonio, Giuliana Palumbo, Lucia Rizzica. 2016. Incentivi E Selezione Nel Pubblico Impiego (Incentives and Selection in Public Employment). *SSRN Electronic Journal* . [[Crossref](#)]
71. Robert J. Stonebraker, Gary S. Stone. 2015. Too Old to Teach? The Effect of Age on College and University Professors. *Research in Higher Education* **56**:8, 793-812. [[Crossref](#)]
72. Cory Koedel, Kata Mihaly, Jonah E. Rockoff. 2015. Value-added modeling: A review. *Economics of Education Review* **47**, 180-195. [[Crossref](#)]
73. Dan Goldhaber. 2015. Exploring the Potential of Value-Added Performance Measures to Affect the Quality of the Teacher Workforce. *Educational Researcher* **44**:2, 87-95. [[Crossref](#)]
74. Lucia Rizzica. 2015. The Use of Fixed-Term Contracts and the (Adverse) Selection of Public Sector Workers. *SSRN Electronic Journal* . [[Crossref](#)]
75. Raicho Bojilov. 2014. Estimating the Effects of Incentives When Workers Learn about Their Ability. *SSRN Electronic Journal* . [[Crossref](#)]
76. Peter Hinrichs. 2014. What Kind of Teachers are Schools Looking for? Evidence from a Randomized Field Experiment. *SSRN Electronic Journal* **25**. . [[Crossref](#)]